# THE NEW MODERN MEDICINE

## Disease, Evidence, and Epidemiological Medicine



JONATHAN FULLER

# The New Modern Medicine

# The New Modern Medicine

## *Disease, Evidence, and Epidemiological Medicine*

JONATHAN FULLER

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries.

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America.

© Oxford University Press 2025

This is an open access publication, available online and distributed under the terms of a Creative
Commons Attribution-Non Commercial-No Derivatives 4.0 International licence (CC BY-NC-ND 4.0),
a copy of which is available at https://creativecommons.org/licenses/by-nc-nd/4.0/.
Subject to this license, all rights are reserved.




Inquiries concerning reproduction outside the scope of the above should be sent
to the Rights Department, Oxford University Press, at the address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer.

CIP data is on file at the Library of Congress

ISBN 9780190066147

DOI: 10.1093/9780190066178.001.0001

Printed by Marquis Book Printing, Canada

The manufacturer's authorized representative in the EU for product safety is
Oxford University Press España S.A. of Parque Empresarial San Fernando de Henares,
Avenida de Castilla, 2 – 28830 Madrid (www.oup.es/en or product.safety@oup.com).
OUP España S.A. also acts as importer into Spain of products made by the manufacturer.

# Contents

*Preface and Acknowledgments* vii

Introduction 1

PART 1. THE NEW MODERN MEDICINE

1. The New Medical Model 21

PART 2. DISEASE

2. Theories of Contagion and Cancer 71
3. Causation and Classification of Disease 117
4. The Nature of Chronic Diseases and Mental Disorders 163

PART 3. EVIDENCE

5. Medical Evidence and Prediction 215
6. Medical Interventions and Risk 266
7. Problems of Extrapolation 307
8. Therapeutic Skepticism 356

PART 4. EPIDEMIOLOGICAL MEDICINE

9. Twenty-First-Century Medicine: Epidemiological, Personalized, or Postmodern? 399

*Appendix: Effect Measures and the Sufficient-Component Cause Model* 437
*References* 443
*Index* 481

# Preface and Acknowledgments

The name of the degree emblazoned on my diploma from my undergraduate studies reads "Bachelor of Medical Sciences, Honors Specialization in Physiology." Having mostly taken courses in the "basic medical sciences" during the program, I graduated with the impression that the science of medicine was laboratory science and that clinical medicine was something like applied physiology. Introductory books on the history of medicine often reinforced this idea, sometimes identifying the rise of "modern medicine" in the nineteenth or twentieth century with transformative discoveries in laboratory research or with "laboratory medicine" or "biomedicine." In medical school, there were times when this afterimage was further burned into my retinas, especially during preclinical training in the classroom and the dissecting room, when I rehashed much of what I'd previously learned. But I came to see that this picture was at best incomplete and left out the impressive role of epidemiological sciences in molding our understanding of disease and evidence, our medical concepts, and our medical reasoning.

It was the philosophy of medicine that first clued me in to the missing story of epidemiology in modern medicine, of the epidemiological medicine, or new modern medicine (as I call it), that doctors and patients now inhabit. This book is an attempt to puzzle through that story. It is foremost a work of philosophy of medicine and philosophy of science, the disciplines in which I trained and now do research. To the extent possible in a single book, I engage with history, medicine, and health science research too, and I hope that historians, health providers, and health scientists will find something of value in it.

I am grateful to the National Institutes of Health for proving funding for this project through a National Library of Medicine Grant for Scholarly Works in Biomedicine and Health. I am also grateful to the University of Johannesburg for funding through a Research Associateship with the Centre for Philosophy of Epidemiology, Medicine, and Public Health that allowed me to make this book open access.

I am indebted to the late historian of medicine Michael Bliss for encouraging me to write a book on the new modern medicine, and to the

philosopher of medicine Ross Upshur for first inspiring me to think about contemporary diseases and medical evidence and to pursue a life in philosophy of medicine. I am grateful to Lucy Randall, my editor at Oxford University Press, for supporting the book and for her patience. I thank the many helpful colleagues who provided feedback (often extensive feedback) on drafts of one or more of the book's chapters, including Robert Aronowitz, Robyn Bluhm, Alex Broadbent, Nancy Cartwright, Benjamin Chin-Yee, Olaf Damman, Marina DiMarco, Benjamin Djulbegovic, Hamed Tabatabae Ghomi, Fabian Hundertmark, Ina Jäntgen, Jason Karlawish, Donal Khoswrowi, Elselijn Kingma, Lucie Laplane, Mathew Mercuri, Chris Phillips, Anya Plutynski, Lauren Ross, Miriam Solomon, Daniel Steel, Jacob Stegenga, James Tabery, David Teira, Steven Tresker, Ross Upshur, Sean Valles, Somogy Varga, Peter Zachar, and anonymous reviewers for Oxford University Press. My sincere apologies to anyone I have inadvertently omitted. I also thank attendees at a book workshop hosted by the Center for Philosophy of Science at the University of Pittsburgh, a reading group attended by members of the Department of History and Philosophy of Science at the University of Cambridge, and students in my Philosophy of Medicine graduate seminar.

Many others provided feedback on articles containing ideas that made their way into this book, and I am grateful to them too. Those articles are:

**Fuller, J.** 2024. "Demarcating Scientific Medicine." *Studies in History and Philosophy of Science* 106: 177–185.

**Fuller, J.** 2022. "Preventive and Curative Medical Interventions." *Synthese* 200(2): 61. https://doi.org/10.1007/s11229-022-03579-0.

**Fuller, J.** 2021. "The Myth and Fallacy of Simple Extrapolation in Medicine." *Synthese* 198: 2919–2939.

**Fuller, J.** 2020. "Epidemiologic Evidence: Use at Your 'Own Risk'?" *Philosophy of Science* 87: 1119–1129.

**Fuller, J.** 2020. "*Medical Nihilism* by Jacob Stegenga: What is the Right Dose?" *Studies in History and Philosophy of Biological and Biomedical Sciences* (currently *Studies in History and Philosophy of Science*) 81: 1–3. https://doi.org/10.1016/j.shpsc.2020.101270.

**Fuller, J.** 2018. "Meta-Research Evidence for Evaluating Therapies." *Philosophy of Science* 85: 767—780.

**Fuller, J.** 2018. "Universal Etiology, Multifactorial Diseases, and the Constitutive Model of Disease Classification." *Studies in History and*

*Philosophy of Biological and Biomedical Sciences* (currently *Studies in History and Philosophy of Science*) 67: 8–15.

**Fuller, J.** 2018. "What Are Chronic Diseases?" *Synthese* 195: 3197–3220.

**Fuller, J.,** and L. J. Flores. 2015. "The Risk GP Model: The Standard Model of Prediction in Medicine." *Studies in History and Philosophy of Biological and Biomedical Sciences* (currently *Studies in History and Philosophy of Science*) 54: 49–61.

There may be small areas of overlap in text, and I thank the original publishers for permission to republish this text.

Last, I thank my family and friends for their enduring support, despite their understandable puzzlement about what it means to undertake research in the philosophy of medicine.

# Introduction

> It seems scarcely credible that in so many directions in so short a time the entire outlook of the science of medicine can have been so revolutionized.
>
> —William Osler (1907)


## Summary

Contemporary scientific medicine is a new modern medicine, one shaped by new disease epidemics, new norms of evidence, and the new sciences of epidemiology. As a contribution to the philosophy of medicine, this book provides a philosophical examination of the new modern medicine and its new medical model, epidemiological medicine. The aim of the book is to provide a picture of the new modern medicine, mostly through a chapter-by-chapter study of its pieces. In this introductory chapter, I compare diabetes and diabetes care in the early 1900s to diabetes today to draw out contrasts between epidemiological medicine and the model of scientific medicine that preceded it, laboratory medicine. I introduce seven problems of the new modern medicine to motivate the project: molecular reductionism; multifactorial etiology; chronicity, asymptomaticity, and variable pathology; evidence hierarchies; meaning of medical risk; external validity of clinical trials; and biased evidence.


## I.1. The Making and Remaking of Modern Medicine

It is 1874. A newly trained physician named William Osler is attending to one of his first patients, an English visitor to Montreal. The man has smallpox. His disease's course will be quick and deadly. Unable to cure the patient's sickness, unable to do little more than offer morphine to ease the pain,

*The New Modern Medicine*. Jonathan Fuller, Oxford University Press. © Oxford University Press 2025.
DOI: 10.1093/9780190066178.003.0001

Osler engages in a ritual practiced by countless doctors before him: he waits at the bedside with his patient, waits for the disease to take its course. The son of a clergyman, Osler delivers one last kindness and reads the Christian Commendatory Prayer at the man's passing (Bliss 2011).

The late 1800s are a time of great transition for medicine. After thousands of years of humoral theory and bloodletting and a reluctant submission to disease, Western medicine is becoming scientific, and according to the historian Michael Bliss (2011, 6), William Osler is the "[h]igh priest of the coming of modern medicine." Osler, a transitionary figure, emerged from the fog of earlier nineteenth-century medicine and represents in the minds of many the meeting of bedside medicine with science. I adapted the title of this section from Bliss's (2011) book, *The Making of Modern Medicine*, which begins with the scene I just described, ten years before the worst smallpox outbreak in Montreal's history. The second case history that Bliss provides "[a]s a guide to the coming of modern medicine" (4) is the founding of the Johns Hopkins medical school in 1893 by Osler and others. Hopkins became the model for learning the ways of scientific medicine, where medical students combined scientific education in the lecture theater and the laboratory with patient care on the hospital wards. Bliss's third case history is the discovery of insulin through physiological research at the University of Toronto in 1922, one of the first realizations of the therapeutic potential of the laboratory model of disease and medical care. Modern medicine had arrived, syringe in hand.

With the epigraph that begins this chapter and our study, Osler marvels at all the change taking place in the science of medicine. The quote is from his introduction to a textbook titled *Modern Medicine*, edited by Osler and Thomas McCrae and published in 1907. But his words would apply equally well if they had been written a hundred years later. If the story of modern medicine has a long fourth act, it goes something like this. Of the many children with type 1 diabetes who were saved by insulin, Ted Ryder was the first to live for seventy years (Bliss 2011). By the time Ryder died in 1993, medicine was the same, yet somehow different.

Insulin does not cure diabetes; it prolongs life with diabetes. And so diabetes became one of the many prolific chronic diseases of twentieth-century medicine, which by the end of the century was confronted with something of a "diabetes epidemic." In the early years of the insulin era, diabetes pathology (deficiency of insulin-secreting cells) as well as primary diabetes treatment (insulin replacement) was conceptually simple. But by the turn of the twenty-first century, diabetes was recognized as a complex and variable

disease of multiple subtypes with multifactorial etiology, including genetics and other risk factors. While physiology and biochemistry were central to the understanding of diabetes in the early 1900s, in the early 2000s diabetes is equally a disease of protein ligands, receptors, and molecular signaling pathways. And while diabetes was once recognized only when it was symptomatic, it is now often diagnosed when it is silent, before it has inundated the body with its harmful manifestations.

Contemporary treatment for diabetes does not aim to cure but instead to forestall the progression of the disease and prevent its complications, including other noncommunicable conditions like heart disease. The watchword for contemporary diabetes management is "risk reduction"—reducing the risk of visual deficits, kidney damage, and heart disease. Rather than physiological first principles or case studies, current "evidence-based" management of diabetes is generally "based" on epidemiological evidence about populations, even while it is supposed to be "individualized" for the patient. Although much of this evidence is "gold standard" evidence from randomized clinical trials, evidence on diabetes therapies has not escaped concerns about the taint of bias, including bias related to industry sponsorship of research.[1] Diabetes care is contemporary scientific medicine in all its intricacy. Understanding scientific medicine and the philosophical problems under its skin will be our task.

This book is about modern medicine, but it is not about Osler's modern medicine. It is about the *new* modern medicine. It is about contemporary diseases, their etiology, and their nature. It is about evidence-based medicine, its evidence, and its inferences. It is about contemporary scientific medicine and its new medical model. In short, it is about medicine around the turn of the twenty-first century and its distinctive problems. It provides a philosopher's view of these problems, for they are, in part, philosophical problems. Before we get to the meat of the philosophy, this chapter will introduce the book and provide a taste of its subject, the new modern medicine. In brief, contemporary scientific medicine is a new modern medicine, one shaped by new disease epidemics, new norms of evidence, and the new sciences of epidemiology.

[1] For example, the drug rosiglitazone came to market in 1999, but its approval was withdrawn in several international markets after a meta-analysis of trials found that it increases the risk of heart attack and death from cardiac causes, a risk that its manufacturer, GlaxoSmithKline, apparently knew about but did not publish or publicly disclose (Nissen 2010).

In section I.2, I use transitions in diabetes care over the past hundred years to illustrate dimensions of the new modern medicine and its new medical model, and I discuss the overall aim of the book: to provide a (partial) philosophical analysis of this new modern medicine. In section I.3, I present seven problems of the new modern medicine to further motivate my analysis. In section I.4, I situate the present work in relation to various scholarly fields. In section I.5, I provide a synopsis of the rest of the book.

## I.2. Portrait of a New Modern Medicine

William Osler died three years before insulin was discovered. However, the model of medicine that gave rise to insulin-dependent diabetes—that is, diabetes kept at bay by insulin injections—was already taking shape, if not fully formed, in his textbook *Modern Medicine*. I will begin this section by using *Modern Medicine*'s chapter on diabetes to illustrate elements of this earlier model of modern medicine. I will compare them to elements of a new model for a new modern medicine, found in contemporary canons of diabetes care.

Diabetes mellitus (from the Greek words for "passing through" and "sweet") is an ancient disease. It was described in the works of Hippocrates, characterized by its telltale sign of frequent evacuation (polyuria) of urine sweetened with (as it turns out) glucose that saturated the blood, usually due to the absence of pancreatic insulin to regulate its metabolism. The disease typically presented in childhood and took the lives of children in months to a few years after their presentation with sickly sweet urine. In the early 1900s, children wasted away as their bodies were drained of calorie-rich sugar, sometimes on diabetes wards like the one at the Toronto General Hospital, where in 1922 exogenous insulin extracted from the pancreas of dogs was injected into its first benefactors and restored them to life (Bliss 2013).

The chapter on diabetes mellitus in Osler's 1907 tome begins by defining the disease through its underlying biological dysfunction and symptoms: "A disorder of carbohydrate and fat metabolism characterized usually by progressive loss in weight, thirst and polyuria, and by the persistent excretion of glucose in the urine while the individual is on a diet containing only moderate amounts of carbohydrates, or, in certain instances, even when no carbohydrates are ingested" (Futcher 1907, 747). The author does not distinguish among etiological subtypes of diabetes due to an incomplete picture of the causal basis of the disease. They do however suspect heredity as well as

pancreatic lesions as contributing to most cases of diabetes. Support for the role of pancreatic pathology comes from physiological experiments in which the pancreas is removed from animals and diabetes follows. The etiology of the disease is viewed through the theoretical lens of digestive physiology as well as carbohydrate biochemistry.

The author notes that the disease "runs a very rapid course," leading to death by diabetic coma in children and very young people (Futcher 1907, 787). The recommended treatments are mainly for individuals manifesting symptoms and consist in dietary restriction of carbohydrate intake, opium, as well as treatment of diabetic complications, include diabetic coma (through alkalis) and gangrenous infection (through amputation). Treatment recommendations are justified mainly through the experience of the author or other named authorities, or through an understanding of the biology of the disease. On treatment with pancreatic extracts, the author writes, "Pancreatic preparations have been given a fair trial, but the results have been disappointing. . . . Pancreatic juice and extract have failed to produce any material benefit" (796). They attribute treatment failure to the lack of any diagnostic test that could distinguish diabetes due to pancreatic disease from diabetes due to other causes, believing that pancreatic extract is a specific treatment only for those cases of diabetes that are due to a specific cause, namely, a deficiency of an unknown pancreatic "ferment."[2]

In the medical model represented in this work, which we can call "laboratory medicine" (Jewson 1976), disease was understood through the lens of biological theories or explanations furnished by experimental laboratory sciences, especially physiology, biochemistry, and bacteriology. Different types of disease were thought to be distinguishable through their different specific causes. Diseases were generally symptomatic and were in principle acute, or short-lived; they often ended in death or natural remission but could potentially be cured or reversed (if not at present). The goals of removal of disease by eliminating a harmful agent or reversal of disease by restoring normal physiology were held up high by the model. Finally, it promoted reasoning therapeutically through pathophysiological thinking and through the experience of clinicians and scientists experimenting with potential remedies in the lab or at the patient bedside (though clinicians like Osler were often nihilists about most of the treatments available to them).

[2] Insulin, contained within whole pancreatic extract, lowers blood sugar in diabetes regardless of the diabetes subtype. Although no information about the means of preparation and delivery of the pancreatic extract is given in this chapter, treatment failure is more likely due to the low insulin

Today the kind of diabetes resulting from destruction of insulin-producing cells in childhood is known as "type 1 diabetes." The irony of the success of insulin is that by allowing those suffering from type 1 diabetes to live, in some ways insulin transformed the disease into a chronic condition and increased its burden for the individual and for society (Feudtner 2003). Still, type 1 diabetes makes up a small slice of the prevalence of diabetes worldwide. The more common type is type 2 diabetes, which usually presents in adulthood and does not initially require insulin. In all, medicine is now confronted with a diabetes epidemic, and patients with diabetes often have multiple other epidemic noncommunicable diseases, including cardiovascular diseases and cancers. I will now compare the laboratory medicine model found in Osler's textbook with the model of diabetes found in a recent edition of *Harrison's Principles of Internal Medicine*, a popular contemporary medical textbook.

The chapters on diabetes mellitus in *Harrison's* (Powers 2015a, 2015b, 2015c) begin by characterizing diabetes as follows: "Diabetes mellitus (DM) refers to a group of common metabolic disorders that share the phenotype of hyperglycemia. Several distinct types of DM are caused by a complex interaction of genetics and environmental factors" (2015b, 2399). Hyperglycemia rather than symptoms define the disease. In fact, the author notes that a long asymptomatic period often precedes the detection of type 2 diabetes. Accordingly, the author outlines asymptomatic screening as a major component of diagnosis. Several etiological subtypes of diabetes are recognized, as well as "pre-diabetes," in which one is at "increased risk for diabetes" (2399). The statement above attributing causation to a complex interaction of genetics and environmental factors is a classic description of "multifactorial" causation. A table lists several risk factors for type 2 diabetes, while two sections describe the "genetics" (rather than "heredity") of diabetes. Type 1 diabetes is understood to be a disease of insulin deficiency due to destruction of insulin secretory cells in the pancreas, while type 2 diabetes is due to complex and variable mechanisms of insulin deficiency and insulin resistance. While physiology and biochemistry contribute to the discussion of the pathogenesis and pathophysiology of diabetes, molecular biology now features prominently: insulin's synthesis, secretion, and action on target membrane

concentration in the extract. Insulin is produced by endocrine cells that comprise a minority of pancreatic cells. Most of the rest of the cells of the pancreas produce digestive enzyme, including enzyme that degrades proteins like insulin. Treatment of diabetics with pancreatic extract generally failed before researchers at the University of Toronto devised physiological experiments that isolated insulin-producing cells in dogs as well as biochemical techniques for extracting insulin from the resulting pancreatic extract (Bliss 2013).

receptors is explained through protein structure and molecular pathways in the cell.

The principles for "management" of type 2 diabetes are listed as "individualized glycemic control," "treatment of associated conditions" (dyslipidemia, hypertension, obesity, coronary heart disease), and "screening for/managing complications of diabetes" (including cardiovascular disease). The author notes that most diabetes treatment focuses on preventing or reducing the risk of complications rather than eliminating symptoms. There remains no cure for diabetes. Details of diabetes management are borrowed from the American Diabetes Association ([ADA] 2021a, 2021b, 2021c, 2021d), which produces regularly updated evidence-based guidelines providing "standards of medical care in diabetes." Because guidelines play a larger role than textbooks in directing medical management today, I will turn to the source.

In a passage that epitomizes the new medical model's approach to treatment, the ADA writes, "Diabetes is a *complex*, *chronic* illness requiring continuous medical care with *multifactorial risk-reduction* strategies beyond glycemic control. Ongoing diabetes self-management education and support are critical to *preventing* acute complications and *reducing the risk* of long-term *complications*. Significant *evidence* exists that supports a range of interventions to improve diabetes outcomes" (2021b, S1, my emphasis). Targeted blood glucose control is the central pillar of treatment and is typically achieved using exogenous insulin or oral antidiabetic agents. The guidelines also advise the treatment of hypertension and the initiation of a cholesterol-lowering statin for those over forty years of age or with atherosclerotic cardiovascular disease. Recommendations are graded using the ADA's evidence grading system, which ranks evidence from well-conducted randomized trials highest (grade A), and clinical experience lowest. All the interventions above are supported by grade A evidence.

In this new medical model, disease is now understood through the additional lens of epidemiological sciences and molecular biological theory. While etiologic subtypes of disease are often identified, disease is also thought to have multifactorial etiology, including genetic causes. It is accepted that disease is often asymptomatic or silent, as well as chronic or enduring. The goal of preventing, treating, and managing diseases and their manifestations by reducing "risk" (as opposed to removing or reversing disease) is central. Finally, in treatment, rather than pathophysiologic reasoning and clinical experience, the model privileges inferences from evidence provided

by epidemiologic studies, especially randomized controlled trials (RCTs). (Clinicians practicing in this model are still skeptical about the effectiveness of many treatments, often because of concerns that the epidemiologic evidence is biased.)

I call the new medical model "epidemiological medicine." Its defining feature is that clinical medicine is infused with epidemiological science. It is "new" in comparison to the old "medical model," laboratory medicine, in which clinical medicine is more closely aligned with laboratory sciences. By the early 1900s, laboratory scientific medicine was beginning to dominate orthodox Western medicine as an ideal. It would assume an unmatched global presence and mark an era of modern scientific medicine—Osler's modern medicine. However, over the twentieth century, medicine underwent a threefold epidemiologic transition in which new disease epidemics, new sources of epidemiologic evidence, and new epidemiological sciences transformed medicine, as illustrated by the transitions in diabetes and diabetes care. The result was the rise of a new medical model and a *new modern medicine*, the contours of which this book will explore.

In this book, I will investigate many of the elements of the new modern medicine identified above: molecular explanations, multifactorial causation, silent chronic diseases, extrapolations and other inferences from epidemiological evidence, risk and risk reduction, and study bias and therapeutic skepticism. I will offer a picture of the new modern medicine in the form of a puzzle. The chapters in Parts 2 and 3 of the book examine the individual pieces of the puzzle just listed, with Part 2 focusing on disease and Part 3 on evidence. Meanwhile, Parts 1 and 4 examine the overall picture and the thread that pulls the puzzle pieces together, epidemiological medicine.

A picture is a partial and imperfect representation of reality. Likewise, many puzzles are missing pieces or have pieces that don't fit together nicely. This book is not an exception. It leaves much puzzling for others to do. No doubt other pictures can be drawn and have been drawn of medicine. We shouldn't insist on hanging just one. While I specifically sketch a picture of contemporary scientific medicine in the early twenty-first century, some philosophers have offered pictures of medicine generally (Broadbent 2019) or epistemic movements in contemporary scientific medicine (Solomon 2015). In works too numerous to list, other philosophers have focused on disease or specific diseases or on evidence or evidence-based medicine (EBM). Historians and other scholars have drawn their own pictures.

For decades, the leading picture of contemporary medicine has perhaps been the "biomedical model" or "biomedicine," which is sometimes referred to simply as "the medical model." The biomedical model serves as a target for scholarly, medical, and patient critics throwing darts at mainstream medicine, lambasting, among other things, its tendency to "reduce" all things in its path to the biological. This picture was drawn by the psychiatrist George Engel (1977) in the 1970s. "The dominant model of disease today is biomedical," he wrote (130). That picture is now faded and missing key details. While some have tried to fit recent developments in scientific medicine into the biomedical model's picture frame—for instance, by accusing EBM or contemporary epidemiology of being "reductionist" (Miles 2009; Krieger 2011)—the picture itself needs updating. Medicine is ruled by a new medical model, epidemiological medicine. This book attempts to piece it together as well as understand it piece by piece.

## I.3. Seven Problems of the New Modern Medicine

Many medical commentators have noted a curious phenomenon. Compared to the middle of the century, by the end of the twentieth century the public appeared more disillusioned with mainstream medicine despite its considerably fuller therapeutic armamentarium.[3] Over the past several decades, medical insiders and outsiders have written about the practical problems of medicine; including problems of biological complexity, multiple etiology, chronic pathology, hierarchical conceptions of evidence, impersonal care, unrepresentative trials, and unreliable evidence. Some of these problems are unique to the new modern medicine, while others have simply become more prominent.

Here, I briefly introduce seven problems of the new modern medicine. They will resurface in chapters to come in the order in which they are presented here.

1. *Molecular reductionism.* Our understanding of disease is often reductionist, explaining the occurrence of disease through sub-personal mechanistic explanations, especially explanations that privilege genetic

[3] Several popular histories have commented on this paradox, including Porter (1997), Le Fanu (2012) and Burnham (2015). Matthew Zack et al. (2004) found that health-related quality of life and self-rated health among U.S. adults generally decreased between 1993 and 2001.

and molecular causes. Not only do reductive molecular explanations exclude psychological, environmental, and social causes, but they may fail to capture the complexity of contemporary diseases that involve large networks of causes operating at multiple nested levels of organization.

2. *Multifactorial etiology*. While acute bacterial infections, the paradigm diseases of the old medical model, are often described as having a single specific and necessary (bacterial) etiology—"*the* cause of the disease"—noncommunicable diseases are typically said to be "multifactorial" in their etiology, caused by a constellation of risk factors. The lack of specific etiology complicates noncommunicable disease prevention not only by multiplying etiological complexity but also because no single cause is necessary for disease occurrence.
3. *Chronicity, asymptomaticity, and variable pathology*. Many contemporary diseases, including mental disorders, are chronic. Chronicity can result in an accumulation of misery and the accumulation of further chronic conditions. However, chronic diseases can sometimes also remain silent and asymptomatic, raising the specter of medicalization. Last, chronic diseases and mental disorders are often variable in their pathology, which creates challenges for explaining, predicting, and treating them using targeted therapies.
4. *Evidence hierarchies*. EBM "hierarchies of evidence" rank the quality of medical evidence. Due to their focus on epidemiological evidence, relevant kinds of non-epidemiological evidence are often ignored, while the reliability of epidemiological evidence such as RCT evidence can be overstated. Moreover, in recommending the single best piece of evidence, hierarchies neglect the importance of bringing together multiple pieces of evidence—each with its own unique relevance—to bear on the medical hypothesis at hand.
5. *Meaning of medical risk*. Risk-related concepts dominate medical dialogue and decision-making. There are challenges of understanding the relationship to and relevance for individual patients of population-level measures of risk and risk reduction, as well as the concept of the patient's risk and their risk reduction. Misunderstandings can lead to suboptimal decision-making for patients and clinicians struggling to make meaning of epidemiological evidence for patient care.
6. *External validity of clinical trials*. The external validity of a clinical trial refers to the extent to which the trial results predict the results of

intervening elsewhere. Because of strict eligibility criteria, many trials are not representative of actual clinical practice and thus may not have good external validity. How to cross the gulf between study populations and real-world populations is a troubling and largely unanswered question, but given the importance now placed on trials, it is one in need of urgent attention.

7. *Biased evidence*. Meta-research has documented many problems with evidence from clinical epidemiology, including biases in publishing trials, biases created by industry funding of trials, and the nonreplicability of many studies. Given these meta-research results, as well as the conflicts of interest that plague medical research, we should wonder to what extent we can trust medical evidence and how we can decide how effective and safe medical interventions truly are.

These seven problems of the new modern medicine are problems *for* medicine. Many of them also arise *from* the new medical model, due to its concepts, goals, inferences, and assumptions.

They are not always the focus of the discussion in the rest of the book but will help to motivate the discussion. They are key leads and witnesses to problems that are worthy of forensic investigation. They are by no means the only problems that plague the new modern medicine. Not included are other important problems that are discussed by physicians, philosophers, and other scholars, including the stalled pace of therapeutic innovation, disillusionment among doctors, the proliferation of the worried well, the popularity of alternative medicines, the spiraling cost of healthcare (Le Fanu 2012); overmedicalization and objectification of patients (Varga 2024); healthcare rationing, managed care, health system errors, and iatrogenesis (Burnham 2015). While these other problems are largely social and systems problems, the seven problems I described above are especially relevant to the study of epidemiological medicine because they result at least partly from the influence of epidemiology on medicine.

I will return to the seven problems throughout the book as they are raised in discussion. I will give them a kind of treatment with which medicine is relatively unfamiliar—a philosophical treatment. I won't pretend that philosophy can fully resolve them from the armchair. But philosophy can contribute to our understanding of them and to our understanding of how they might be resolved. For example, whether there exist epidemics of molecularly based, multifactorial, and silent chronic diseases depends partly on the

world but also, partly on the ways we view, carve up, and define contemporary diseases. On the other hand, the problems of evidence hierarchies, unrepresentative trials, medical risk, and biased epidemiologic studies are special problems for certain kinds or principles of evidence. The practical problems of medicine are entangled with philosophical problems. Philosophers can be of service by disentangling the philosophical from the practical problems and in so doing can contribute toward unraveling both.

## I.4. A Contribution to the Philosophy of Medicine

This book is foremost a work in the philosophy of medicine. I understand the philosophy of medicine broadly as a field of philosophy investigating any aspect of health, the provision of healthcare, and research on health, using any philosophical approach (Broadbent and Fuller 2020).[4] That said, my focus in this book is necessarily narrower. I write mainly of disease, medical evidence, and the science of medicine.[5] I also treat problems that are mainly metaphysical and epistemic rather than ethical, not because ethical problems aren't intertwined with these other problems (they are) nor because they are unimportant (they are not) but because ethical problems have (still, after decades of work in philosophy of medicine) received far greater attention than other philosophical problems in medicine.

This book is also a work in the philosophy of science in that I write about scientific medicine, its theory, its methods and tools, and its inferences. Some chapters adopt a history and philosophy of science (HPS) approach by integrating history into the philosophy of science. While this book can also be seen as a contribution to the history of science and medicine, it is a philosophical contribution, providing a synthetic account of contemporary scientific medicine and its history, and drawing only selectively on the historical literature. In a humble philosophy monograph, I cannot do justice to the vast literatures in history of science and medicine nor in sociology of medicine,

[4] There have been a number of edited volumes in philosophy of medicine in recent years, including Solomon et al. (2017) and Broadbent (2025). The philosophy of medicine has often been defined more narrowly so as to exclude bioethics or philosophy of psychiatry, or as a subdiscipline of the philosophy of science (Gifford 2011; Solomon et al. 2017). Arguably, the time for restrictive definitions has passed (Broadbent and Fuller 2020).

[5] Likewise, my approach is mainly analytic, using methods such as analysis and argument, and naturalistic, engaging with work in medicine and medical science (in addition to philosophy).

anthropology of medicine, nor other scholarly disciplines studying science or medicine.

One of the main functions of the history in this book (including in this chapter) is to give what we might call "contrastive focus" to the philosophy. By contrasting scientific medicine roughly around 2000 (and later) with scientific medicine roughly around 1900 (and earlier), I draw out several philosophically (and historically) unique problems and features of the new modern medicine.[6] Though I aspire to provide a unifying account of contemporary scientific medicine, this contrast allows me to selectively examine elements that are distinctive to contemporary scientific medicine, such as those stemming from the epidemiologic transitions to which I alluded in section I.2. Even still, I cannot treat nearly all distinctive elements of contemporary scientific medicine in one book. For example, I mostly set aside psychiatry to limit my scope (except in Chapter 3, in which the relevance of my theory of chronic disease for mental disorders is too great to ignore). Many of the problems of contemporary psychiatry are shared with the rest of contemporary medicine, but others are special. There is also a question about whether psychiatry is or should be part of contemporary medicine and its analysis, the very question that led Engel (1977) to write about the biomedical model. It is a question that is better left until after contemporary medicine is analyzed.

This book is more about the theory rather than the practice of medicine. I do not provide data about the way that practicing physicians understand disease or the way they reason from evidence. The book is about our new medical model, about how scientific physicians are supposed to think about disease and evidence in the ideal. It is divided into several parts, with Parts 2 and 3 comprising most of the chapters. Part 2 is about disease, especially theories of disease, while Part 3 is about medical evidence, especially inferences from evidence. Disease and evidence are perhaps the two biggest topics in philosophy of medicine, those that have been most studied, which justifies giving each topic its own part in the book. Much of the existing literature on disease in philosophy of medicine is about concepts of disease (and related health states), and much of the literature on evidence is about EBM, its hierarchy of evidence, and the reliability of RCT evidence and evidence

[6] This idea of contrastive focus is inspired by accounts of contrastive causal explanation that use a contrast to foreground a particular cause of a phenomenon while backgrounding other causes (Lipton 2004).

of biological mechanisms. In this book, I shift the discussion to less explored topics.

Several historians and doctors focus on the effectiveness of contemporary medical interventions or on new social roles for contemporary medical interventions as the key to the coming of medicine modern (Wootton 2007; Le Fanu 2012; Greene et al. 2016). Contemporary treatment could have been a third major topic for this book alongside disease and evidence. However, I will argue in Chapter 1 that effective medicine was secondary to scientific medicine; scientific medicine precipitated effective medicine and better characterizes developments in medicine during the 1800s and 1900s. Moreover, medical interventions are indeed studied in this book, particularly in Part 2 when I discuss therapeutic prediction, preventive and curative medical interventions, and therapeutic skepticism.

Last, emerging/re-emerging infectious disease epidemics (epidemics of new or evolving infections) could have been another topic treated in this book, one that describes a characteristic contemporary threat to health (Fauci 2022). Alongside noncommunicable disease epidemics of cancer, cardiovascular disease, and chronic diseases, emerging/re-emerging infectious diseases are the new epidemics. I do examine infectious disease in Chapters 2 and 3 and use examples of emerging/re-emerging infections such as tuberculosis, HIV infection, and COVID-19 throughout. I do not investigate epidemics per se because this book is about medicine rather than public health, and epidemics, as a population-level phenomenon consisting in an increased population burden of disease, are foremost a public health problem (see Fuller (2022a) on explaining and intervening in epidemics). However, in Chapter 9, I use the COVID-19 pandemic as a case study for assessing the new modern medicine's prognosis.

These caveats aside, I will now provide a glimpse of the philosophical discussion to come.

## I.5. Synopsis

In Part 1 of this book, I discuss the genesis of the new modern medicine and its new medical model, epidemiological medicine (Chapter 1). In Part 2, I investigate disease, including theories of contagion and cancer (Chapter 2), causation and classification of disease (Chapter 3), and the nature of chronic diseases and mental disorders (Chapter 4). In

Part 3, I investigate evidence, including medical evidence and prediction (Chapter 5), medical interventions and risk (Chapter 6), problems of extrapolation (Chapter 7), and therapeutic skepticism (Chapter 8). In Part 4, I return to epidemiological medicine and compare it to competing visions of contemporary medicine: personalized/precision medicine and postmodern medicine (Chapter 9).

In Part 1, Chapter 1, I defend the view that modern medicine begins with scientific medicine against the view that it begins with effective medical interventions. Then, rather than arguing that scientific medicine is a science, I propose that scientific medicine is best understood as the integration of clinical medicine with modern research sciences. Not only does this alternative capture what made medicine modern by around the turn of the twentieth century, but it also foregrounds the finding that different sciences competed for the mantle "the science of medicine" in the 1800s and 1900s. By the early 1900s, laboratory medicine was becoming established as the dominant model of scientific medicine as medicine modeled itself after experimental laboratory sciences. But over the 1900s, medicine underwent a threefold epidemiologic transition in disease, evidence, and science, resulting in a new medical model I call "epidemiological medicine" that is modeled after epidemiological sciences, and a new modern medicine in which epidemiology enjoys paramount clinical priority.

One recuring theme from Part 2 of the book is that there are (perhaps surprisingly) theories of disease in contemporary medicine. These theories include empirically supported theories of families of disease, such as a theory of contagious diseases (germ theory) and one of cancers (a gene-centric somatic mutation theory), examined in Chapter 2. One of their main characteristics is that they integrate and unify a family of diseases that share a common etiological pattern. Contemporary medical theories also include more conceptually oriented theories of disease etiology, including the sufficient-component cause theory discussed in Chapter 3, which represents the multifactorial etiology of diseases, including noncommunicable diseases. It has been suggested that infections adhere to an older model, a theoretically and practically useful "monocausal model" of disease classification that imparts a specific universal etiology. I argue that the monocausal model is a myth: specific pathology rather than specific etiology was the key conceptual innovation of modern medicine. Finally, when medicine lacks an explicit theory of disease, as it presently does for its heterogeneous collection of chronic diseases, metaphysics can potentially supply one. In this case, a

dispositional theory can help illuminate the nature of chronic diseases—as well as mental disorders—by understanding them as incurable dispositions that persist even when their characteristic symptoms and manifestations are silent (Chapter 4). That these chronic dispositions may lack specific pathology helps to clarify why they pose a particular challenge for medical explanation, prognosis, and intervention. My scientifically and theoretically motivated approach to disease in Part 2 is a needed counterweight for a philosophy of medicine that historically approached the problem of understanding disease mainly through traditional conceptual analysis.

A theme that emerges in Part 3 is that by paying more attention to the relationship between evidence and inference, EBM could overcome many of its current challenges, including those that require integrating diverse evidence. In Chapter 5, I reconstruct the implicit view of evidence in EBM in which medical evidence consists in measurements of risks (including "risky" measures of the effect size), which are extrapolated from epidemiological studies and applied to individual patients. I contrast it with an argument theory or inference theory in which medical evidence supports premises in arguments for or against medical hypotheses. In Chapter 6, I investigate how and in what sense preventive and curative medical interventions reduce the risk of medical conditions in populations. I also probe the concept of the individual patient's risk in medicine, arguing that it is best interpreted as a real-world property but that—given important problems with deriving patient properties from epidemiological evidence—medicine should reinterpret these probabilities in an inferential light as representing epidemiological uncertainty. In Chapter 7, I argue that extrapolating from clinical trials and other epidemiological studies faces two distinct problems of extrapolation, one of constructing extrapolation arguments and one of determining what evidence can support those arguments. I examine solutions to the former problem that involve dissecting mechanisms and causes (including my own solution, mechanism-cofactor analysis), and I propose a new solution to the latter problem that harnesses structural analogies. Finally, in Chapter 8, I scrutinize Medical Nihilism, the view that we should have little confidence in the effectiveness of medical interventions, advancing three objections. As an alternative response to the problem of biased evidence, I propose Meta-Evidentialism, which holds that we should have reduced confidence in the accuracy of therapeutic evidence as a measure of therapeutic effects and argues that EBM should include meta-research evidence in its therapeutic

evaluations. Meta-Evidentialism views bias as primarily a problem of faulty inferences rather than faulty study results.

In Part 4, Chapter 9, I revisit epidemiological medicine and the seven problems of the new modern medicine. I show how these seven problems result at least partly from the influence of epidemiology on medicine. I then consider the possibility that my picture of the new modern medicine is already outdated or soon will be, replaced by some new vision of scientific medicine such as a personalized/precision medicine modeled after laboratory molecular biology or by a postmodern medicine in which scientific medicine has lost its authority. However, precision medicine is epidemiological medicine with a molecular gloss, while epidemiological medicine remains in reasonably good health, despite the persistence of pseudoscientific medicines like homeopathy and resistance to scientific medicine during the recent COVID-19 pandemic. In fact, these examples reveal directions in which epidemiological medicine is expanding, while also exposing its weaknesses and limits. I conclude that the picture of the new modern medicine sketched in this book remains a faithful (if incomplete) depiction of scientific medicine today and provides philosophical diagnoses for many of its ailments.

## I.6. Conclusion

Contemporary scientific medicine is a new modern medicine, shaped by new disease epidemics, new norms of evidence, and the new sciences of epidemiology. Its new medical model is an epidemiological medicine. This book provides a picture of the new modern medicine and its new medical model, motivated by seven problems: molecular reductionism; multifactorial etiology; chronicity, asymptomaticity, and variable pathology; evidence hierarchies; meaning of medical risk; external validity of clinical trials; and biased evidence. As a contribution to the philosophy of medicine, it provides a philosophical treatment of contemporary scientific medicine and philosophical diagnoses for many of its problems.

# PART 1

# THE NEW MODERN MEDICINE

# 1
# The New Medical Model

## Summary

Orthodox Western medicine underwent a historic transformation in the 1800s and 1900s, becoming the modern medicine of today. At the same time, medicine today seems different compared to medicine a century ago—a new modern medicine. What made medicine modern? And what changed? The short answer is that modern medical science made medicine modern, and the newfound influence of epidemiological sciences on medicine in the second half of the 1900s led to a new model of scientific medicine, epidemiological medicine.

I start with the first question, arguing against the narrow view that effective medicines made medicine modern. Rather, what fundamentally changed was that medical interventions and the rest of medicine were newly shaped by modern research sciences. Scientific medicine made medicine modern. In modern scientific medicine, clinical medicine is integrated with and modeled after modern research sciences. In the early twentieth century, these sciences were primarily experimental laboratory sciences, and the dominant model of scientific medicine in society—the medical model—was laboratory medicine. Turning to the second question, over the twentieth century medicine underwent a threefold epidemiologic transition in which its diseases, evidence, and science were transformed by new epidemics and new epidemiological sciences with new research methods. This transformation led to a new medical model after the Second World War, epidemiological medicine. While there is something to the story that laboratory medicine, biomedical science, and biomedicine are representative of research priorities in the postwar era, it overlooks the clinical priority that epidemiological science and epidemiological medicine enjoy in the new modern medicine.

*The New Modern Medicine*. Jonathan Fuller, Oxford University Press. © Oxford University Press 2025.
DOI: 10.1093/9780190066178.003.0002

## 1.1. Medicine's Transformation

Around the turn of the nineteenth century, a patient with fever attended by a reputable orthodox physician in North America and parts of Europe could expect to be asked to recount their experience, including their fever's course, duration, and quality. They might be told that their sickness was the result of tension in their arteries or an overabundance of a bodily humor (in this case, blood). If there was an ongoing epidemic of many people succumbing to a similar fever, its upstream cause might be miasmas—toxic vapors rising from the ground. And the patient might be treated with bleeding to reduce the tension in their arteries or to restore their body's balance of humors.

By the early twentieth century, that same patient might instead provide samples of bodily fluids for culturing, microscopy, and a few laboratory tests. They might be diagnosed as infected with a germ, perhaps one of a known kind. And they might be given new treatments developed in the laboratory and trialed at the bedside, such as arsphenamines or serum therapy, alongside older interventions, such as quinine and mercury, that remained in the physician's medicine bag.

In the early twenty-first century, the patient's encounter might again go differently. They might undergo more extensive diagnostic testing, including molecular diagnostics, to cover a greater number of known pathogens as well as other causes of fever—now more common—such as cancer or chronic inflammatory conditions. The physician might also consider the probability of a specific diagnosis, the risk of disease complications, and the quantitative effectiveness of potential treatments—which by now include myriad treatments vetted in randomized clinical trials.

Something clearly changed. But what?

Recognizing the dramatic nature of the changes medicine underwent in the nineteenth and twentieth centuries, some historians, doctors, and other writers have referred to this period as the rise of "modern medicine." In the *Birth of the Clinic*, Michel Foucault wrote, "Modern medicine has fixed its own date of birth as being in the last years of the eighteenth century" (1973, xii), in the aftermath of the French Revolution, when medical education and practice in Paris reorganized itself around the teaching hospital, *la clinique*, and doctors began to see disease with a "modern medical gaze" that pierced through the skin to the underlying anatomical lesion. Others disagree on modern medicine's exact birthdate and circumstances, while agreeing that medicine became modern around this time or soon after. For instance,

David Wootton (2007) locates the birth of modern medicine in the year 1865 in Glasgow. It was there that Joseph Lister applied the germ theory of putrefaction to the problem of surgical infection and invented the method of surgical antisepsis, the sterilization of surgical sites and surgical instruments with a germicidal agent. Others don't refer to "modern medicine" but still recognize that orthodox Western medicine underwent dizzying changes in the 1800s and 1900s.

While Foucault emphasizes a new way that modern doctors regarded patients and their diseases, Wootton emphasizes new effective interventions that modern doctors had available. What both innovations have in common is the role that science played in their development. The North American physician William Osler is often seen as the very embodiment of modern medicine of the late nineteenth century, marrying bedside medicine to medical science (Bliss 2011). Hearkening to this personification, Osler described "[m]odern scientific medicine" as the place where "two great types meet": "the scientific and the practical physician" (1908, 189).

However, there are many sciences for the practical physician to meet. Which of them ground scientific medicine? In *The Evolution of Modern Medicine*, Osler ([1913] 2013) associated the rise of modern scientific medicine with the anatomical and laboratory sciences of the nineteenth century. Laboratory sciences are central to what has been called "laboratory medicine," "biomedicine," or—projecting a sense of its authority—"the medical model," in which the fruits of these sciences are applied to clinical medicine. We saw in the introduction that in several important respects the laboratory medicine model is outdated as a rendering of contemporary medicine. Over the twentieth century, medicine underwent a threefold epidemiologic transition with the rise of new disease epidemics, new forms and standards of evidence, and the new sciences of epidemiology. The outcome of medicine's continuing transformation was a new medical model, an epidemiological medicine rooted in epidemiological science. This book explores the new modern medicine, its diseases, its evidence, and its new medical model.

Modern medical science made medicine modern, and epidemiology especially shaped the new model of scientific medicine, epidemiological medicine. In section 1.2, I reject the view that modern medicine began with modern effective medicines; modern medicine instead began with scientific medicine. In section 1.3, I discuss different ways of demarcating modern scientific medicine from nonscientific or pseudoscientific medicine, suggesting that what is distinctive about modern scientific medicine is that

it is integrated with and modeled after modern research sciences. In section 1.4, I distinguish different models of scientific medicine and discuss features of the previous medical model, laboratory medicine. In section 1.5, I summarize the threefold epidemiologic transition of the twentieth century that gave rise to the new medical model of today, epidemiological medicine. In section 1.6, I introduce the elements of epidemiological medicine and argue that it has greater clinical priority today compared to laboratory medicine and better characterizes the new scientific medicine than does the biomedical model.

## 1.2. Modern Medicines?

Unlike art history or intellectual history, the history of Western medicine is not divided neatly into agreed-upon premodern and modern periods. But recognizing the continuity between this period in medical history and present-day medicine, historians and doctors have sometimes suggested that in the nineteenth century, or by the twentieth century (at least in Western Europe and North America), medicine was becoming modern, if it was not modern already.[1] Of course, calling this period "modern" as opposed to something else is a choice. But it is a choice meant to signal something dramatically different about this new medicine. The question of what made medicine modern is really, then, about the nature of medicine's transformation in the 1800s and 1900s.

A popular idea I will examine in this section is that the development of modern medicines—the drugs, devices, and surgeries of today—was what made medicine modern. I will examine variations on this suggestion in which modern medicines are those that are uniquely effective and find them all to be inadequate or insufficient for capturing fundamental changes in medicine over this period. While historians, philosophers, and doctors are

[1] For example, Foucault (1973), Shryock (1974), Wootton (2007), Bliss (2011), Le Fanu (2012), Burnham (2015). I will not consider all previous and potential narratives of modern medicine in this chapter. Roger Cooter (2011) surveys various meanings of medical modernity. In addition to its association with science (especially experimentation, mechanism, and the laboratory), we can also understand modern medicine through the bureaucratization of medicine and medical institutions in the tradition of the sociologist Max Weber (1948), the objectification of the body in the tradition of Foucault (1973) or N. D. Jewson (1976), or through industrialization in medicine and society.

ultimately free to define modern medicine however they like, there are good reasons to avoid an effectiveness criterion.

It is sometimes said by doctors, scientists, and historians that the coming of today's treatments signaled the arrival of modern medicine. In *The Rise and Fall of Modern Medicine* (2012), James Le Fanu locates "the rise of modern medicine" in the immediate post–World War II period with the "pharmaceutical revolution" that brought a cornucopia of new medical technologies and especially drugs. The subtitle of Bliss's (2011) *The Making of Modern Medicine* is *Turning Points in the Treatment of Disease*. The book recounts three case studies that Bliss calls "turning points in humanity's assault on disease" (4), beginning with a story of therapeutic despair in the setting of a smallpox epidemic in 1885 and culminating in therapeutic triumph with the discovery of insulin in 1922, a span of time representing the "coming of age of modern medicine" (1). And David Wootton (2007) argues that modern medicine was born in 1865 with the discovery of surgical antisepsis, which prevented deadly postoperative infections and allowed surgery to soar to new heights. According to Wootton, the transition to modern medicine was completed in 1941 with human trials of medicine's first antibiotic, penicillin. In a phrase, modern medicines made medicine modern (Greene et al. 2016).

A natural question then arises: What makes these medicines modern? What distinguishes them from the use of mercury, bloodletting, and other medicaments in earlier times? The answer cannot reside in the nature of mercury and bloodletting. Mercury is a metal. Yet iron, lithium, and other metals are part of modern medicine's armamentarium, while mercury is not. And bloodletting is used today for a limited range of diseases, usually those resulting from excess iron in the blood. It must be the way that bloodletting is used today that makes contemporary bloodletting, or phlebotomy, modern and late eighteenth-century bloodletting not modern. If we wish for "modern medicine" and "modern medicines" to mean something more than current medical care and therapeutic practice, if we want these terms to mark a longer period in medical history characterized by a dramatic shift from what came before, we must say more about what it means for medicine and its medicines to be modern. What Le Fanu, Bliss, Wootton, and others have in mind when they refer to antibiotics, insulin, or surgical antisepsis as uniquely modern is that they are uniquely effective for the diseases against which they are used. In comparison, mercury and bloodletting were generally ineffective as they were used in the 1700s and earlier—and even harmful.

Charles Rosenberg (1977) rejects the notion that earlier therapeutics such as bloodletting, diuretics, and emetics were ineffective, and in so doing raises trouble for those who wish to use effectiveness to mark the beginning of modern medicine. Rosenberg notes that the therapeutic systems like humorism within which these therapies were used were effective in several senses: they produced visible effects like blood loss, they produced a sense of understanding between doctor and patient, and they provided the patient reassurance. Thus, perhaps a notion of effectiveness that is appropriately sensitive to the meaning of therapeutics within their historical context cannot be used to mark the beginning of modern therapy and modern medicine.

Yet those who argue that effective therapeutics and modern medicine begin with therapeutic revolutions in the 1800s or 1900s usually have different senses of effectiveness in mind: effectiveness as the ability to cure disease and save lives (Wootton 2007; Bliss 2011; Le Fanu 2012). According to Wootton, there was very little medical cure before 1865. Hippocratic medicine, much like modern medicine, aimed to cure disease; judged against this standard, it failed. In comparison, modern medicine succeeds, at least somewhat.

As Wootton confesses, Wootton's (2007) story is "Whiggish" in that it adopts a historical narrative of progress (Broadbent 2019). However, my main misgiving about Wootton's curative effectiveness criterion (modern medicine begins with the discovery of effective curative therapies) is not that it is Whiggish but that it is insufficient for capturing fundamental changes in medicine over the 1800s and 1900s. The curative effectiveness criterion fails to emphasize the transformation in therapeutics that allowed therapy to become effective. In other words, it focuses on a consequence of medicine's modern transformation—its newfound curative success—rather than the nature of this transformation. Moreover, it too narrowly focuses on therapeutics and downplays other developments that are part of the very same transformation. As a result, the criterion leads Wootton to dismiss historical changes predating Joseph Lister's 1865 discovery of surgical antisepsis that are rightly considered part of medicine's modernization.

One difficulty facing Wootton's dating of the beginning of the transition to modern medicine to 1865 is that it comes roughly seventy years after Edward Jenner's discovery of smallpox vaccination. While smallpox vaccination prevents rather than cures smallpox, because it saves lives Wootton

cannot in good faith look past its effectiveness. Nor can Wootton look past the importance of this discovery in light of the historical importance of smallpox as a cause of death. To deal with the problem of smallpox vaccination, Wootton argues that this discovery, unlike the invention of surgical antisepsis, did not involve the application of scientific understanding (and consequently the discovery of smallpox vaccination was not replicated for other infections until Pasteur wielded the developing science of bacteriology to develop new vaccines in the second half of the 1800s). In contrast, Lister's discovery involved applying a scientific theory, the germ theory, which held that surgical infections (and infections more broadly) were caused by living microscopic germs.

Following Wootton's logic, we can say that what was modern about the invention of surgical antisepsis was the role played by medical science.[2] However, if we accept that what makes a therapeutic innovation such as surgical antisepsis modern is the application of science, it is unclear why we should also require that the innovation is effective. Medical science—even good science—often fails to deliver effective interventions but is not necessarily less scientific in its failing. Wootton understandably wishes to draw attention to the incomparable curative effectiveness of modern medicine. But if we understand modern medicine primarily through its association with science, then its effectiveness is evidence that science has the potential to generate effective interventions (along with harmful ones) rather than implied by the concept of modern medicine.

Moreover, the curative effectiveness criterion overlooks an equally important application of science to therapeutics: the use of systematic clinical testing to determine which treatments are effective and—just as important—which treatments are ineffective or harmful. Progress in therapy depends as much on discarding ineffective and harmful treatments as on discovering effective ones (Shryock 1974). Arguably, what makes medicine scientific and modern is not that surgical antisepsis works against infection while bloodletting does not, but that medicine has determined through scientific research which one works and which one doesn't. Lister kept track of mortality rates among his patients before and after he introduced antisepsis into his surgical practice (Wootton 2007). But earlier in the century, before Lister, doctors

[2] A similar story can be told about the discovery of insulin, which is used by Bliss (2011, 2013) to mark the transformation to modern medicine and involved applying scientific method and theory from disciplines like physiology and biochemistry.

doctoring in the tradition of Paris clinical medicine made empirical testing a routine part of medicine. Practicing in large hospitals that housed many inpatients, disciples of Pierre Louis and his "numerical method" evaluated treatments like bloodletting by counting and comparing outcomes in groups of hospital patients who were following different therapeutic regimens (Matthews 1995),[3] often finding their effectiveness unimpressive. It is reasonable to think that Paris doctors practiced modern medicine because they used a scientific method to study medical practice—and did so in the absence of new and effective therapies.

A defender of curative effectiveness could modify the criterion such that modern medicine is defined by overall progress in curative effectiveness rather than by the discovery of effective cures. Then, to the extent that discarding ineffective cures constitutes progress, Paris clinical medicine was progressive and modern. However, this modification overlooks further developments in Paris medicine that are also part of medicine's scientific transformation and that are important for *eventual* therapeutic progress yet do not in themselves constitute improvements in the overall effectiveness of medicine.

Wootton places great importance on clinical trials for determining which treatments are effective. He thus commends Paris clinical medicine for counting and comparing patients using consistent case definitions, which was central to the numerical method. In the next breath, Wootton downplays the importance of a new way of viewing the body advanced by Paris clinicians, which Foucault (1973) earlier argued marks the birth of modern medicine. Paris clinicians adhering to the "pathological anatomy" of Xavier Bichat distinguished different kinds of disease according to their characteristic anatomical lesions. They learned to correlate anatomically defined diseases, observed postmortem using the naked eye on the autopsy table, with diagnostic signs elicited premortem on physical examination at the patient's bedside. This innovation permitted more reliable diagnosis and prognosis. Pulmonary tuberculosis, as a disease of tubercular lesions in the lungs, could be diagnosed in the living by listening with the stethoscope to a

[3] British "numerists" did so even earlier, in the eighteenth century (Tröhler 2005, 2011), including through population studies on the safety of variolation for smallpox prevention (Bird 2019). Variolation involved inoculating people with smallpox pus to give them protection against full-blown smallpox. As variolation was a precursor to smallpox vaccination, the scientific research on variolation makes it more difficult for Wootton to overlook the relevance of smallpox vaccines for modern medicine.

patient's abnormal breath sounds. The patient's (often grim) clinical course could then be forecast more reliably owing to similarities in the clinical trajectory among patients with tubercular lung lesions. Paris pathological anatomy involved the application of scientific knowledge to disease classification, diagnosis, and prognosis, especially knowledge from the science of morbid anatomy. It is thus part of medicine's scientific maturation. But because pathological anatomy had diagnostic and prognostic relevance rather than immediate therapeutic relevance, Wootton argues that it is *not* part of modern medicine.

Wootton's conclusion is misguided for several reasons. First, the innovations instilled by pathological anatomy facilitated therapeutic progress in two ways. On the one hand, pathological anatomy provided a scientific basis for consistent case definitions (Matthews 1995), which (as Wootton recognizes) allowed Paris clinicians to reliably count disease outcomes in populations and make inferences about whether therapies are effective using the numerical method. Pathological anatomy could thus help medicine become more effective overall by discarding worthless treatments.

Pathological anatomy also initiated a "specificity revolution" (Rosenberg 2002) in medicine that was crucial for the eventual development of effective remedies. By tying disease classification and diagnosis to specific kinds of pathology, pathological anatomy created a fruitful template that was later exploited by nineteenth-century physiologists as well as by bacteriologists who defined diseases based on infection with a specific bacterium. Tuberculosis, now understood in the West as infection with *Mycobacterium tuberculosis*, can be cured with an antibiotic treatment that targets its specific pathology. Therefore, by permitting consistent case definitions and by grounding disease kinds in specific pathology, pathological anatomy facilitated the therapeutic progress that Wootton celebrates. Although these innovations resulted from the direct application of science to diagnosis rather than to therapy, they are part of the broader transformation that allowed medicine to become more effective.

Beyond all its consequences, the nineteenth-century scientific transformation in diagnosis and prognosis is fundamental in its own right because diagnosis and prognosis are fundamental to clinical medicine. Broadbent (2019) argues that throughout history, medicine has always been engaged in the project of inquiring into diseases through explanation and prediction for the purpose of curing diseases—medicine is at its core an inquiry. In contrast,

Wootton (2007) argues that medicine is fundamentally a technology for treating illness rather than a science or an intellectual pursuit, which is why he characterizes modern medicine in terms of therapeutic developments.[4]

However, even if we accept that medicine is a technology rather than an inquiry, it is unclear why we should see it only as a technology for treating illness rather than as additionally a diagnostic and prognostic technology, a tool for finding out about and forecasting disease. Many of the historic developments in medical technology over the nineteenth century were in diagnosis, from the development of the stethoscope in the early 1800s to the use of quantitative "precision instruments" such as the calibrated thermometer, the microscope, and X-rays in the late 1800s. The latter technologies were closely identified with a new "scientific medicine" (King 1991; Matthews 1995).

Similarly, even if cure or treatment of disease is the ultimate goal of medicine (Wootton 2007; Broadbent 2019), diagnosis and prognosis are core activities and intermediate medical goals, if not ultimate goals. Diagnosis and prognosis greatly occupied doctors practicing premodern orthodox medicine just as they now occupy doctors practicing modern medicine, as a major component of what doctors do. To properly characterize medicine in terms of what it does and not only in terms of its highest aspiration, we mustn't overlook diagnosis and prognosis. All three pursuits—diagnosis, prognosis, and intervention—were transformed with modern medicine, a point that the curative effectiveness criterion obscures.

As I will argue further in the next section, what made medicine modern was not its newly effective therapies but its new contract with science in all dimensions, including in therapy, but also diagnosis and prognosis. In fact, in modern medicine, effective therapy often depends on scientific diagnosis because modern treatment is often tailored to specific, pathologically understood, consistently defined diseases. Even in therapy, medicine's modernization is best understood not through acquiring effective treatments nor by becoming progressively more effective overall—these were upshots of medicine's modern transformation. Rather, what makes modern medicines modern is their relationship to science: *scientific therapy* rather than effective therapy. In short, modern medicine is scientific medicine. How we should

[4] As Broadbent points out, this leads Wootton to wrongly diminish the very inquiry that led to effective therapies like invasive surgeries in the post-antisepsis era, which required the accumulation of anatomical knowledge over preceding centuries.

best understand "scientific medicine" and carve its historical boundaries is the topic of the next section.

## 1.3. Modern Scientific Medicine

Many historians and doctors have identified the birth of modern medicine with the rise of scientific medicine (Cooter 2011).[5] Scientific medicine is an ambiguous concept, and in this section I consider two ways of characterizing it that implicate different "demarcation problems": the problem of demarcating science from non-science, and the problems of demarcating science-based medicine from medicine not grounded in science. To help us understand modern medicine, a successful solution to scientific medicine's demarcation problem must distinguish scientific medicine from earlier as well as contemporaneous medical sects or traditions, a task for this section. To help us understand the new modern medicine, a solution must also be able to accommodate changes in scientific medicine over time, a problem I approach in the next section.

The challenge of understanding what makes medicine modern can be phrased as a problem of demarcation: What demarcates modern medicine from other medical traditions? What made medicine modern in the nineteenth or twentieth century? What distinguished the modern medical tradition from other traditions at the time? And—even more important for the theme of this book—what makes medicine today still modern and yet different, a *new* modern medicine? My starting suggestion is that the establishment of scientific medicine as the dominant medical tradition is what makes medicine modern. Thus the problem becomes partly one of demarcating scientific medicine from nonscientific medicines synchronically (at a given time) and diachronically (over time). In answering this challenge, we should also clarify the relationships among scientific medicine, orthodox Western medicine, and heterodox medicines.

The North American context is a good place for resolving these problems because of the prominence of heterodox medical sects and traditions in

[5] For examples of historians identifying modern medicine or the new medicine of the 1800s/1900s with science or scientific medicine, see Rothstein (1972); Foucault (1973); Shryock (1974); Fye (1987); King (1991); Warner (1995); Porter (1997); Marks (1997); Bynum (2006); Bynum et al. (2006); Quirke and Gaudilliere (2008); Knight (2009); Bliss (2011); Ackerknecht (2016); Hruban and Linder (2022); Roth (2022).

the 1800s and 1900s and because the transition to scientific medicine in America has been well studied (Rothstein 1972; Fye 1987; King 1991). The history of orthodox Western medicine is usually traced from Ancient Greek Hippocratic medicine in the fifth century BCE to the nineteenth century, if not later, and can be identified over time through its embrace of the Hippocratic tradition, including Hippocratic writings, Hippocratic ideas, and Hippocratic therapies (Conrad et al. 1995). In nineteenth-century America, the orthodox tradition was often known as "regular medicine" and was contrasted with "irregular medicines" such as Thomsonianism, hydropathy, and homeopathy. John Harley Warner (1997) argues that the strongest identity marker for regular medicine over the nineteenth century was its therapeutic practices and especially its therapeutic principles, including allegiance to bloodletting and other "heroic" treatments. Therefore its allegiance to traditional Hippocratic therapies allows us to identify the orthodox Western medical tradition over time despite the revolutionary changes it underwent, and to say that orthodox Western medicine *became* scientific medicine. The lancet-worshiping orthodox doctors started out the century committed to traditional medical systems, but by the end of the century they were increasingly committed to science. In fact, science gradually replaced the lancet as the main symbol and identifier of orthodox Western medicine (Warner 1997).

There is an abundant literature in philosophy of science—both classic and more recent—on "the demarcation problem": What demarcates science from non-science, and especially from pseudoscience (Popper 1963; Pigliucci and Boudry 2013; Gordin 2021)? Thus I will first examine whether scientific medicine should be understood as a science and whether the problem of demarcating scientific medicine from earlier and contemporaneous medicines should be understood as one of picking out what made orthodox medicine a science by the late nineteenth or early twentieth century when these other medicines did not become science (and were even plausibly pseudoscience). According to this suggestion, orthodox Western medicine became a science with the transition from traditional Hippocratic medicine to scientific medicine, while homeopathy and other heterodox medicines never transitioned to science. There were lively debates in the 1800s and 1900s among leading medical figures about whether orthodox medicine was, could be, or should be a science, rather than, for instance, an "art" (King 1991; Matthews 1995; Lawrence and Weisz 1998). Speaking to la Société Médicale d'Observation, Louis proclaimed that "medicine is a science" (1837, 3). Experimental physiologist

Claude Bernard later agreed, writing, "Medicine is a science and not an art" (1947, 50).

Likewise, philosophers of science have recently asked what demarcates medicine as a science[6] from other legitimate knowledge practices and from pseudosciences. Applying a general solution to the demarcation problem in science developed by Paul Hoyningen-Huene (2013), Alexander Bird (2019) and Somogy Varga (2021) argue that its systematicity in various dimensions distinguishes contemporary (orthodox Western) medicine as a science from other activities, including commonsense knowledge. Varga argues that its greater systematicity also demarcates medicine, as a science, from homeopathy, as a pseudoscience (though Naomi Oreskes [2019] argues that systematicity is an insufficient demarcation criterion). Meanwhile, Bird argues that clinical medicine became systematic starting in the eighteenth century by assessing the effectiveness of interventions like variolation for smallpox prevention using statistical comparisons, a forerunner to Louis's therapeutic application of the numerical method.

However, modern scientific medicine is most associated with the numerical method and pathological anatomy early in the nineteenth century, as well as with the research laboratory and medical technologies imported from the laboratory beginning later in the nineteenth century and continuing into the twentieth century (Matthews 1995).[7] The numerical method, pathological anatomy, and laboratory research were themselves either research sciences or scientific methodologies, as mainstream doctors understand science today. The numerical method, as a rudimentary method of quantifying outcomes, establishing correlations, and making population-level clinical comparisons, can be seen as a precursor to contemporary clinical epidemiology (Morabia 2006). Pathological anatomy combined the science of morbid anatomy with the method of "clinicopathological correlation," correlating diagnostic and prognostic clinical signs with postmortem anatomical findings. And

[6] A previous generation of philosophers provided reasons why we should *not* consider medicine a science (e.g., Munson 1981).

[7] Jeremy Greene and colleagues observe, "Whether medicine has changed by revolution or evolution, a substantial corpus of medical historiography since Osler has continued to emphasize the role of science in making medicine modern" (2016, 4). Greene and colleagues criticize the historical narrative of modern scientific medicine because they suggest that it construes medical history as necessarily progressive. However, we need not view the relationship between science and medicine as progressive or unproblematic. Historians have argued that the immediate impact of science or the rhetoric of science on nineteenth-century orthodox medicine was to elevate the profession rather than to benefit patients (Shortt 1983; Warner 1995). In the introduction, I identified seven problems of the new modern medicine. As we will see in Chapter 9 and throughout this book, these problems are partly created by medicine's relationship with medical science, particularly epidemiology.

laboratory research consisted in sciences like experimental physiology, biochemistry, and bacteriology. Thus the development of scientific medicine in the nineteenth century is primarily characterized by a tightening relationship between clinical medicine and scientific research, *by clinical medicine becoming science-based* rather than by medicine becoming an independent science of its own (Knight 2009; Roth 2022).[8]

Osler compared William Harvey with Thomas Sydenham to make a similar point. While both were (seventeenth-century) physicians, Harvey's greatest contributions to medicine were in scientific research and theory (he discovered, through physiological experiment, the circulation of the blood), while Sydenham's greatest contributions were arguably in medical practice (he wrote treatises on the diagnosis and treatment of diseases and elevated practice above theory). Writing at the end of the nineteenth century, Osler declares, "Harvey and Sydenham, types of the scientific and the practical physician, though contemporaries, were uninfluenced, so far as we know, by each other's work or method. . . . Modern scientific medicine, in which these two great types meet, had its rise in France in the early days of this century" (1908, 189). Osler helped found the Johns Hopkins University medical school in 1893, which became a model for training physicians in scientific medicine in America (Bliss 2011; Hruban and Linder 2022).

My solution to scientific medicine's demarcation problem relies on the common distinction between the category of medical research on the one hand and that of clinical medicine or medical practice on the other. The term "medicine" is sometimes applied to one category, sometimes to the other, or to the conglomeration of both. Although a thorough examination of the distinction is beyond our scope, a simple working pair of definitions will allow me to explore the historical and contemporary contours of scientific medicine: medical (scientific) research is the pursuit of general medical (scientific) knowledge, while medical practice is the pursuit of patient care (which might involve producing *particular* scientific knowledge about, for instance, the patient's diagnosis).

The research/practice distinction I have drawn might seem blurred in Paris clinical medicine because pathological anatomy and the numerical method were clinical science (Warner 1995). Paris clinicians conducted autopsies in the hospital morgue, identified clinical signs at the bedside

[8] Whether its close relationship with scientific research was enough to make clinical medicine a science or whether clinical medicine was already a science is not a question I will answer here.

and established their connection to autopsy findings, and compared cases with the same diagnosis on the hospital ward to establish the effects of medical interventions (Ackerknecht 1967). Science was thus carried out in the process of patient care. Likewise, in contemporary scientific medicine, doctors caring for patients enrolled in clinical trials are also involved in both research and practice. However, the distinction I drew between research and practice was at the level of (overlapping) types of activities, not at the level of token activities. Medical research is a type of activity where the goal is general knowledge, while medical practice is a type of activity where the goal is patient care. The same token activity can serve multiple goals at once, so it is no threat to my distinction that Paris clinicians often contributed to medical practice *and* medical research simultaneously in the act of applying the stethoscope to the patient's chest.[9]

Modern scientific medicine, as science-based medicine, did not begin with the arrival of what we now recognize as medical sciences such as anatomy, morbid (pathological) anatomy, and physiology. Modern scientific anatomy is commonly said to have begun in the sixteenth century, modern scientific physiology in the seventeenth century, and modern scientific morbid anatomy in the eighteenth century (Conrad et al. 1995). Accordingly, one might say that modern scientific medicine began around one of these earlier times with the start of modern medical science.[10] But these sciences that would eventually become part of medicine predate science-based medicine (Knight 2009). Until the 1800s most students were not trained in medicine by spending time in an anatomy theater or a morgue doing dissections or in a lab doing experiments. Even after modern anatomy, physiology, and morbid anatomy debunked the theory of the four humors, thanks partly to Harvey's research on the circulation in the seventeenth century, physicians continued to use bloodletting, an intervention originally justified by Hippocratic humoral theory as a corrective to an overabundance of blood, well into the

[9] The Belmont Report (National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research 1978), which helped establish the contemporary ethics of human subjects research, makes a similar distinction between research and practice on the basis of goals. It notes that research and practice are often carried out together. However, the Report is not explicit about whether it refers to activity types or tokens. Moreover, by defining practice as "solely" aiming to improve patient care, it seems to stipulate that the categories of research and practice do not overlap (Bluhm and Borgerson 2018). I do not make any such stipulation.

[10] Keekok Lee (2012) makes this sort of claim in *The Philosophical Foundations of Modern Medicine*, arguing that modern medicine began by the seventeenth century with the coming of modern science.

1800s (though by then its theoretical justification varied among orthodox practitioners; Warner 1997).

In the eighteenth and early nineteenth century, orthodox medicine was instead based on rationalistic medical systems such as humorism and the monistic system of the American physician Benjamin Rush. These systems consisted in a few general principles or laws governing the pathophysiology of diseases, from which explanations and therapeutic rationale for particular diseases and treatments could be inferred. In Rush's (1805) system, all disease was of a febrile kind and resulted from tension in blood vessels, especially arteries, which could be relieved through depletive therapies—bloodletting and laxatives. Systems were not provided to medicine by an empirical knowledge-generating scientific field of research but by an authoritative figure or canon. There was no eighteenth-century field of humoral research, no science of humorism. Rush supported his system in his voluminous writings with sweeping generalizations from scarce observations and with weak analogies (King 1991). Nineteenth-century framers of scientific medicine often contrasted their "scientific," "empirical," or "positive"[11] methods with these "speculative" or "rationalistic" systems that were founded on mere "theory" or "hypothesis" (King 1991; Warner 1998).

Like other heterodox medicines, nineteenth-century homeopathy was similarly based in a system (King 1991; Loudon 2006), one invented by the German physician Samuel Hahnemann and codified in his 1810 canon, *Organon of the Healing Art*. Hahnemann's system was comprised of two main laws or principles (the law of similars and the principle of dilution). Orthodox medical critics were quick to point out the lack of scientific foundations for these principles in the light of contemporary physical science, but also that the empirical literature on the effectiveness of homeopathic cures was flawed, often consisting in uncontrolled case studies of symptomatic remission following homeopathic treatment rather than comparing those treated with homeopathy to those not treated in a like comparison (Fuller 2025). To the extent that there was a field of homeopathic research, it was not scientific research, or at least it was significantly less scientific than the science produced by the numerical method.

Ultimately a more complete argument that the numerical method and experimental physiology were science while the systems of humorism and

[11] See King (1991) on the multiple meanings of medical empiricism in nineteenth-century America.

homeopathy were not (as mainstream medicine understands science today) would require engaging with the first demarcation problem I set aside, the problem of distinguishing science from pseudoscience. Surely the paucity of sound empirical research supporting nineteenth-century medical systems compared, for instance, with the germ theory and theories of endocrine regulation of the nineteenth century, bolsters my claim that medical systems were not scientific fields of research.[12] However, again, my main purpose in this section is to examine science-based medicine, which I argued best characterizes scientific medicine. I will thus focus on what it is for medicine to be science-based rather than what it is for a field of medical research to be science. (I approach the latter problem in Fuller 2024, 2025 and briefly in Chapter 9.) I will assume with the mainstream medical (and philosophical) consensus that the nineteenth-century fields of physiology and bacteriology were sciences, while the nineteenth-century systems of homeopathy and humorism were not.

In claiming that modern medicine is characterized by science-based medicine or scientific medicine, I mean clinical medicine *integrated with* and *modeled after* scientific research. As suggested above, in Paris clinical medicine the *integration* was partly institutional: medical education, research, and practice were co-located in the hospital. The same clinical case could be the recipient of clinical care, an exemplar for clinical teaching, and an object of clinical research. Integration also included the implementation of diagnostic, prognostic, and therapeutic knowledge supplied by scientific research in the care of the next patient.

Paris clinical medicine also *modeled* itself after scientific research in that it borrowed its disease concepts, clinical orientation, and clinical reasoning from pathological anatomy and the numerical method. In the model of Paris clinical medicine, diseases were defined and diagnosed anatomically; concrete kinds of disease—visible at autopsy after the patient's life was extinguished—were the units of medical diagnosis over and above the living individual; and diagnosis correlated physical signs with physical lesions (an application of pathological anatomy's method of clinicopathological correlation) while treatment predictions were based on inferences from population

[12] Massimo Pigliucci (2013) suggests a two-part theoretical-empirical demarcation criterion: a scientific theory must have internal coherence and logic, as well as empirical support. Compared with the theories of bacteriology and experimental physiology, medical systems generally possessed internal coherence but lacked empirical support (Warner 1997). Empirical support should ostensibly be a requirement for theories that make empirical claims.

comparisons, both relying on the numerical method to quantify, correlate, and compare. Because few orthodox treatments had been studied this way, and those that had been studied generally came out looking unimpressive, doctors practicing in this model were often measured, skeptical, or nihilistic about the benefit of medical interventions. In the first half of the 1800s, Paris clinical medicine was mainly the model for academic elites in Paris or Boston.

It might be objected that throughout the history of medicine before the nineteenth century there were many examples of doctors whose practice contained scientific elements (even if the term "scientist" is strictly a nineteenth-century invention; Knight 2009). After all, doctors have always diagnosed, and diagnosis (done well) is arguably a highly systematic and scientific activity (Varga 2021). Moreover, orthodox medicine since Hippocrates has often adopted a naturalistic worldview, which is characteristic of science (Gordin 2021).

In response, it's worth recalling that with the term "scientific medicine" I mean only to single out medicines in which clinical practice is grounded in fields of scientific research such as pathological anatomy or experimental physiology. By "scientific medicine," I *don't* mean every medicine that is scientific, nor am I denying that some past medicines were scientific in certain ways. Of course, there were earlier isolated examples of physicians who integrated research science into their practice in the sense captured by my term "scientific medicine"; Osler (1908, 189) raises the example of the eighteenth-century anatomist-surgeon John Hunter. However, Osler goes on to describe Hunter as a lone voice crying out in the wilderness. Modern medicine is defined by *schools* of scientific medicine, one of which came to dominate medicine around the turn of the twentieth century (as discussed in the next section). In other words, one scientific medicine practitioner does not a modern medicine make.

Last, it could be argued that the grand theorizing and systematizing that gave rise to systems of medicine such as humoral theory could be considered science, even if not supported by data generated by the kinds of empirical methods favored by the Paris "numerists" or the later experimental physiologists; thus schools of medicine integrating these systems into diagnostic and therapeutic practice, including schools that predate the nineteenth century, such as humoral medicine, deserve to be considered scientific medicine. There are, after all, multiple perspectives on what constitutes science or medical science (King 1982; Laudan 1983), some of which give a central place to theorizing and systematizing.

However, recall that modern scientific medicine is a tradition continuous with present mainstream medicine and comparatively discontinuous with medicine before roughly the nineteenth century. Thus the relevant sciences of modern scientific medicine are more narrowly *modern* medical sciences (Knight 2009), academic fields of research that contemporary mainstream medicine would recognize as science. These fields would include anatomy and physiology but not humorism or homeopathic research.[13] Anatomy and physiology are among the sciences that were included in Osler's prototype scientific medical curriculum at Johns Hopkins around the turn of the twentieth century and that are still included in the mainstream medical curriculum today. Nineteenth-century orthodox physicians sometimes used the term "positive" in referring to these fields or their methods, implying a positivist view of science (King 1991; Warner 1998). "Positive medicine" would indeed be an accurate description of modern scientific medicine.

In short, modern scientific medicine is science-based medicine, clinical medicine integrated with and modeled after modern research sciences. We have already seen that there are multiple medical sciences in which medicine could be based, including pathological anatomy and experimental physiology. In the rest of this chapter, I will argue that the rise in influence of different sciences can be used to mark different eras in modern medicine.

## 1.4. The Medical Model: Laboratory Medicine

If Paris clinical medicine was the birth of modern scientific medicine, then laboratory medicine, a distinct school of scientific medicine, achieved its rise. The laboratory medicine model became *the* medical model by the early twentieth century thanks to the rise in experimental laboratory thinking in medical research and education.

For disciples of the Paris school of clinical medicine in the first half of the 1800s, including foreigners from Europe and America who sought out its elite instruction, Paris medicine was science. Yet today Paris clinical medicine is less often remembered under the rubric of "scientific medicine" compared to what has been called "experimental medicine" or "laboratory medicine," a distinct style of medicine that emerged in the second half of the

[13] Although many homeopaths, including Hahnemann, have considered homeopathic research to be science (Dean 2001), mainstream medicine considers it pseudoscience (Mukerji and Ernst 2022).

1800s, especially at scientific institutions in Germany. Historians often identify a uniquely scientific medicine with the ascent of the medical science laboratory and with the use of technologies, theory, and experimental mindset imported from the laboratory in diagnosis and treatment (Fye 1987; Bynum 1994), following a convention set by several nineteenth-century actors (Virchow [1847] 1956; Bernard 1865; Flexner 1910).[14] The science taken to be "the science of medicine" slowly shifted in the second half of the nineteenth century (albeit with some resistance), from the observational science of pathological anatomy and the numerical method to the experimental laboratory science of physiology, biochemistry, and bacteriology.

Warner (1995) notes that the reason why laboratory medicine is often equated with scientific medicine is that laboratory medicine was the vision that was eventually adopted by the entire orthodox medical profession. Rather than viewing scientific medicine narrowly as the association of clinical medicine with laboratory science (or with some monolithic "medical science"), we should understand scientific medicine as medicine's association with particular modern sciences. Doing so reveals that scientific medicine has changed in remarkable ways over time as various sciences have jockeyed to place themselves as the foremost science of medicine. Louis had pronounced that "medicine as a science, rests entirely on observation" (1834, 15). Beginning in the mid-nineteenth century, many experimental physiologists who served as ambassadors for a laboratory-based experimental medicine rejected the observational or "empirical" nature of Louis's numerical method. "Empiricism precedes science," countered Bernard (1875, 702–703).

Bernard's (1865) ideas in his *Introduction to the Study of Experimental Medicine* are frequently recited. Many of them are also found in the earlier writings of the cellular pathologist Rudolph Virchow. The first article in the inaugural 1847 issue of Virchow's journal *Archives of Pathological Anatomy and Physiology and of Clinical Medicine* was titled "Standpoints in Scientific Medicine." In it, Virchow describes a newfound split between the science and the practice of medicine, in which the physiologist does the science in the

[14] "Scientific medicine" was an important actor's category in the nineteenth and early twentieth century, emphasizing the application of science to clinical medicine. Though "scientific medicine" took on varied meanings (and in early twentieth-century America included both preclinical and clinical science), laboratory research was a common factor (Roth 2022). In this chapter, I sharpen and broaden the actor's category to include other nineteenth- and twentieth-century traditions that share with laboratory medicine the application of research science to clinical medicine, while allowing other research sciences to substitute for laboratory science.

lab, while the scientific practitioner applies it in the clinic. The physiologist discovers the causes of anatomical changes (visible on dissection) by controlling the conditions of life through experiment, particularly animal experiment (vivisection). For Virchow, scientific medicine must be familiar with normal physiology as well as pathological physiology, altered or abnormal physiology. Therapy completes the experimentalist's triad by studying the ways in which pathological physiology can be restored to normal physiology. Virchow, who had earlier trained to become an army doctor, saw "pathological physiology as the fortress of scientific medicine on which pathological anatomy and the clinic are only the outworks!" ([1847] 1956, 449).[15]

Experimentalists from Virchow and Bernard to American physiologists argued that medical science should be seen as a discipline distinct from clinical practice and done by dedicated full-time scientists (Fye 1987; Bynum 1994). By the turn of the twentieth century, the research scientist was no longer always a practicing doctor, and medical research became a "kind of separate estate within medicine" (Bynum 1994, 117). The separation brought new problems of integration for scientific medicine not faced by Louis's clinical science, which were partly solved in elite centers by having medical students learn laboratory science from laboratory scientists, by housing laboratory research in faculties of medicine, by instituting teaching laboratories in medical schools and diagnostic laboratories in hospitals as companions to the research laboratory, and by encouraging medical trainees to undertake careers in laboratory science. Laboratory science was also integrated with clinical medicine by providing the latter with biological theory such as the venerated germ theory as well as laboratory technologies. Early twentieth-century technologies of laboratory medicine included diagnostic tests such as bacteriological culture and microscopy, as well as interventions such as vaccines, diphtheria antitoxin, and thyroid extract.

In the laboratory medicine model, clinical medicine was now modeled after laboratory science. Diseases were understood in biological (as opposed to primarily anatomical) terms through theories or explanations developed in experimental laboratory sciences such as physiology, biochemistry, and bacteriology. In other words, diseases were viewed in the clinic just as they were in the laboratory, as an infection with a specific bacterium (as per the

[15] Virchow (1855) later refined his ideas by founding pathological physiology in the microscopical study of cellular pathology and by emphasizing that it was a science independent from physiology and including clinical observation and autopsy alongside laboratory science (Roth 2022).

germ theory) or a specific disturbance in hormone or chemical physiology—often using the same scientific instruments used in laboratory research to diagnose disease in the clinic. Their specific causes also potentially distinguished different types of disease. And diseases were generally regarded as symptomatic and acute, terminating eventually in either death, natural remission, or cure/reversal.[16] Removal of disease by eliminating an intruding pathogen or reversal of disease by restoring normal physiology were the model's ultimate goals.[17] To achieve these ends, therapeutic reasoning was to emulate laboratory reasoning, applying pathophysiological theory and explanations, the results of laboratory experiments, and the trial-and-error experience with treatments gained at the bedside.

Even if there remained reason for skepticism about existing treatments, the model was optimistic about the potential for exploiting laboratory thinking in the discovery of new interventions. In experimental physiology, this promise took the form of what was called "experimental therapeutics," "physiological therapeutics," or "the principle of physiological antagonism": if disease is pathophysiology (a deviation from normal physiology), then successful treatment consists in correcting pathophysiology to restore normal physiology, which can be achieved by using drugs that theoretically and experimentally oppose or "antagonize" the disease (Bartholow 1881), as in exogenous insulin for lowering pathologically raised blood sugar in diabetes mellitus. In bacteriology, the concept of "internal antisepsis" was touted even before dramatically effective cures were found for infectious diseases: if disease is infection with a bacterium, then cure can be attained by killing the bacterium (Bartholow 1876). As Lister had applied germ theory to the problem of preventing surgical infection, internal antisepsis extended the application of germ theory to the problem of curing infection.

[16] Even prominent infectious diseases that often had a chronic time course—syphilis, tuberculosis—were regarded as potentially curable after the bacteriological revolution. Though long-term "degenerative" noncommunicable conditions like angina pectoris were discussed, they were generally regarded as inevitable diseases of aging and largely beyond medicine's powers (D. Jones 2013). And around the turn of the twentieth century, "chronic disease" was not the influential concept or category that it is today (Weisz 2014).

[17] While the germ theory also spurred efforts in personal and public sanitation (Tomes 1998), these were developments in "public hygiene"—what we would now call public health—rather than clinical medicine. And while new vaccines that could *prevent* disease were another gift of the germ theory, many of the vaccines developed in the early period of modern medicine—including vaccines against anthrax, rabies, cholera, and typhoid—had limited local importance for laboratory medicine's European proponents given the rarity of the human diseases they prevented in Western Europe by the time they came around (Wootton 2007).

Paris clinical medicine and laboratory medicine were originally associated with particular nineteenth-century schools in France and Germany (Ackerknecht 1967). Norman Jewson (1976) refers to the bundle of ideas that each school represented as "medical cosmologies." Calling them "cosmologies" appropriately acknowledges that they became medical *world*views, influencing medicine in Western Europe and North America (and eventually, in the case of laboratory medicine, most of the world) as they became detached from their original local and temporal contexts. Alternatively, we could choose to think of each as a different thought style, épistémè, or disciplinary matrix. Studying scientific communities, Ludwik Fleck (1979) developed the concept of a "thought collective" with a characteristic "thought style," using the origins of modern ideas about syphilis in the 1800s and early 1900s as his chief example. Foucault (1970) wrote about an épistémè, an a priori framework that conditions scientific knowledge. Foucault's earlier work on the modern medical gaze of the postrevolutionary Paris clinic is readily understood using this idea. Finally, several philosophers and historians have analyzed episodes in the development of modern medicine through the lens of Thomas Kuhn's (2012) "disciplinary matrixes" or "paradigms" (Gillies 2005; Wootton 2007; Goodwin 2015).

I will instead refer to the three scientific medicines discussed in this chapter (Paris clinical medicine, laboratory medicine, epidemiological medicine) using the less committed terms "model" or "medicine," for several reasons. First, medical thinkers themselves write about different models of medicine and of different medicines, including the "patient-centered care model," the "chronic care model," "evidence-based medicine," "person-centered medicine," and "precision medicine"—not to mention the "biomedical model" and "biomedicine." These ideas represent either the medical status quo or an aspirational vision for medicine. The biomedical model in particular has long been a central motif in medicine and for scholars studying medicine. I hope to connect to these strands of thought. Moreover, for various reasons, Fleck's, Foucault's, and Kuhn's ideas don't adequately capture the kind of structures I have in mind. For instance, the units of radical historical change for Kuhn are "scientific revolutions" or "paradigm shifts," which are changes within a science; however, the shift from Paris clinical medicine to laboratory medicine was a change in the place accorded to different sciences.

Several clarifications are needed with respect to the three medical models I discuss. First, with these models I don't intend to capture a full structure that would constitute a complete framework for medicine or a unified unit of

medical historical change. A model is a partial representation of some phenomenon. I use medical models as a device for capturing certain features of a very complicated phenomenon, scientific medicine, while setting other features aside. Second, they are foremost normative models or ideals for clinical medicine promoted by medical visionaries, through medical education, through medical research, and through the institutions in which all of the above are embedded (though often not as "models," not explicitly, and not in these exact terms).[18] Although the models also partially represent actual medical practice, their influence was never instantaneous, evenly distributed, or total. Nonetheless, their impact on professional practice and identity make them important guides to medical historical change.

The medical historical change in question sees these different models as appearing, then rising, then falling in clinical priority compared to one another. By "clinical priority," I mean that at different times different models and their characteristic sciences have had comparatively greater influence and power in orthodox clinical medicine. In mid-nineteenth-century France, followers of Louis and Bernard debated whether Paris clinical medicine or laboratory medicine represented proper science, before laboratory medicine won out by the century's end and came to enjoy greater priority in French medicine (Matthews 1995).[19] Older models never completely go away. The legacy of Paris clinical medicine endures in many of the rituals of contemporary medicine, including the physical examination, anatomical dissection in medical school, and the medical autopsy. Moreover, each model may continue to have greater purchase for some diseases or in some specialties. Many medical problems are best described in structural terms, and surgeons continue to operate in a world of anatomical lesions. Nonetheless, the appearance and rise of each model helps us take notice of important changes in scientific medicine, even if the models never completely fall away.

Paris clinical medicine may have been the first school of scientific medicine and was a model emulated beyond Paris in Europe and North America. Yet it did not have universal purchase in orthodox medicine. In

[18] In reconstructing them, I've relied on historical research on nineteenth-century scientific medicine, especially Matthews (1995), Warner (1997), and Bynum (1994), as well as authoritative works by historical actors in France, Germany, and America, especially Louis (1834), Virchow ([1847] 1956), Bernard (1865), and Bartholow (1881).

[19] J. Rosser Matthews (1995) identifies a third group who steered a middle course by studying disease physiologically like Bernard but in the clinic with Louis, especially using precision measuring instruments. Carl Wunderlich, who studied the dynamics of fever quantitatively in the hospital using the thermometer, is representative of this third camp.

mid-nineteenth-century America, disciples of the Paris school comprised only a subset (or subsect) of orthodox practitioners alongside orthodox doctors who did not receive a scientific medical education (or a university education at all). Moreover, American orthodox medicine itself competed fiercely with other medical sects, especially the paradigm pseudoscientific medicine homeopathy, in an unregulated, wide-open marketplace (Warner 1997).[20]

It was only with laboratory medicine that a united scientific medicine[21] achieved global societal dominance, superseding previous and contemporaneous scientific and nonscientific medicines in its influence. Laboratory science became the preeminent tradition of academic medical research, orthodox medical education became more widely aligned with academic medicine, and scientific medicine and scientific medical research became more widely supported by the state (Bynum 1994). Laboratory medicine was now more than just a model for scientific medicine; it was *the* model for medicine. The rise of laboratory scientific medicine started first in Western Europe in the second half of the nineteenth century, and then continued in North America in the early years of the next century, especially following the Flexner Report of 1910, gaining momentum as the twentieth century rolled on.

Supported by the recently established Carnegie Foundation, Abraham Flexner toured Canada and the U.S. in the opening years of the 1900s to identify the best (and worst) medical schools. He advocated a formula for medical education that was modeled after Johns Hopkins and German academic medical education. It consisted of two years of learning basic medical sciences (anatomy, histology, embryology, physiology, biochemistry, pharmacology, pathology, and bacteriology) taught by full-time research/teaching faculty, followed by two years of training on the hospital wards. It is a model that most North American schools follow today. For Flexner, this formula for medical training represented a new era, "dominated by the knowledge that

[20] Orthodox American medicine pushed heterodox sects like homeopathy to the fringes through several strategies. After it was formed in the mid-1800s, the orthodox American Medical Association excluded homeopaths and forbade its members from consulting with them (King 1991). Eventually, medicine became more widely regulated, with medical licensure and medical school accreditation favoring orthodox medicine over heterodox medicines (Flexner 1910). As a marker of its privileged political status, orthodox medicine is today considered mainstream medicine, while heterodox schools like homeopathy are considered alternative or complementary medicine.

[21] That is to say the relationship between laboratory science and clinical medicine was increasingly one of cooperation, notwithstanding local episodes of tension, conflict, and hostility due to competing interests, cultures, and epistemologies (Sturdy 2011).

medicine is part and parcel of modern science" (1910, 53). Flexner's vision for scientific medicine was not merely of a medical curriculum that included a healthy amount of scientific knowledge but of one that trained scientific physicians, or "scientist-doctors" (Whitehead 2013), those who applied both scientific knowledge and the scientific method to patient care (King 1982)—in my terms, one who integrated practice with science and modeled practice after science.

For Flexner (1910), the ideal science after which to model clinical medicine was experimental laboratory research rather than the "empirical" science that characterized Paris clinical medicine (Heinbokel 2024).[22] In fact, spending time in the laboratory molded the scientific medical mind. The laboratory and the hospital should be united in spirit (as well as in physical space). As for medical education, "Physiology is, in a sense, the central discipline of the medical school. It is the business of the physician to restore normal functioning: normal functioning is thus his starting point in thought, his goal in action" (63). Not only were diseases to be understood and classified in terms of biological causes and dysfunction, but they were to be treated by reversing dysfunction (the principle of experimental therapeutics). Finally, doctors were to diagnose and treat disease using the mindset of the laboratory scientist by testing hypotheses against the patient's data.

In the years following the Flexner Report, American medical schools perceived to be in nonconformity with the laboratory medicine model shuttered. Meanwhile, around the world, laboratory research continued to produce new technologies for practice. Biochemical assays multiplied, while the physicist's X-rays became commonplace in diagnosis. The screening of chemicals in the laboratory for their antibacterial properties led to the discovery of arsphenamines for syphilis, followed by the sulfa drugs for a range of infections. Insulin was isolated through animal experiments and biochemical techniques. Vitamins were synthesized in the laboratory. And iron lung ventilators and dialysis machines were designed to substitute for the body's failing physiology. Then, in the two decades following the Second World War, the dam holding back therapeutics seemingly broke, as a "pharmaceutical revolution" delivered to medicine a rapid succession of pills and prophylactics, including antibiotics, childhood vaccines, steroids and immunomodulatory drugs, cancer chemotherapy, psychotherapeutics, oral

[22] Flexner (1925) later endorsed a broader vision of science as an attitude rather than as a specific method.

contraceptive pills, and antihypertensive and antidiabetic agents (Le Fanu 2012). The unstoppable growth of the pharmaceutical industry producing these pharmaceuticals was matched by growing public investment in medical research, especially laboratory science, by research institutions like the British Medical Research Council (MRC) and the U.S. National Institutes of Health (NIH).

In short, by the early twentieth century the blueprint guiding medical education and research was one in which clinical medicine was based in experimental laboratory sciences, including bacteriology and physiology. With laboratory medicine becoming established as *the* medical model in the West during the late 1800s and the first half of the 1900s, we can speak of "modern scientific medicine" not only as a kind of medicine practiced, for instance, by earlier practitioners of Paris clinical medicine but as a period in medical history beginning with the rise in dominance of the laboratory medicine model. If this book was the story of modern medicine, it could, like many such stories, end here. But this book is instead about the new modern medicine. In important ways, by the end of the twentieth century laboratory medicine had been overthrown and modern medicine upended by a threefold epidemiologic transition.

## 1.5. Epidemiologic Transitions

In the nineteenth century, acute conditions—infections, poisonings, and injuries—especially preoccupied physicians. The epidemics that wrought death and despair were infectious, spread from animals or other humans (Weindling 1992). In the first decade of the twentieth century, the pages of the *New England Journal of Medicine* were dominated by infectious diseases, and half of the mortality in the U.S. was due to pneumonia, influenza, tuberculosis, and gastrointestinal infections (D. Jones et al. 2012).

That was all to change. Continuing an ongoing trend, developed countries reined in acute infectious diseases over the first half of the century. What came next was described by Adbel Omran (2005) in his much-discussed "theory of the epidemiology of population change." According to Omran, the nineteenth century was an "age of receding pandemics" that gave rise to an "age of degenerative and manmade diseases" in the twentieth century. As mortality from acute infections continued to decline, particularly in early life, people began living longer, initially in developed Western countries, and then in

other parts of the world. Noncommunicable diseases—cardiovascular disease, cancer, and chronic diseases—then replaced infections as the leading causes of morbidity and mortality in these countries as populations aged and were exposed to the harmful effects of a modern lifestyle and environment. At the end of the first decade of the twenty-first century, heart disease rather than infectious disease was the number one killer and along with cancer accounted for two-thirds of all deaths (D. Jones et al. 2012).

Omran described variations on this trend occurring in other countries, particularly developing countries. Others have pointed out many exceptions to Omran's rules of population change, for instance, the arrival of deadly new plagues in the form of pandemics of emerging and re-emerging infectious diseases (Mackenbach 2022), those infections caused by novel pathogens or by older pathogens with increased incidence or geographic range, respectively (Snowden 2008). Emerging infections included HIV/AIDS in the late twentieth century and COVID-19 in the early twenty-first century (the latter a leading cause of death in developed countries during its pandemic phase), while re-emerging infections included those like tuberculosis that had become partially resistant to the antibiotics of the pharmaceutical revolution. Likewise, outbreaks of childhood infections like measles that had been formerly suppressed by the new vaccines crept up in local areas with faltering vaccine coverage (Tiller et al. 2023).

On top of this, as morbidity and mortality from noncommunicable diseases grew, they came to be seen as new pandemics (Weisz 2014). They included epidemics of cancer and heart disease, but also other chronic conditions such as heart failure, chronic obstructive pulmonary disease, diabetes, and even obesity, as populations accumulated cancer-causing mutations, heart disease–causing atherosclerotic plaque, and incurable chronic pathology of all sorts. Moreover, mental disorder diagnoses became more common in terms of both the number of recognized disorders and the number of people recognized as being disordered. In the U.S., an epidemic of mental distress manifesting in rising "deaths of despair" from suicide and substance use after 2000 began to influence and even depress life expectancy, particular for certain social groups (Case and Deaton 2020).

Thus perhaps Omran's "age of receding pandemics" was really followed by an "age of new pandemics." Granted, epidemics are partly created by medical decisions about which conditions to diagnose and at what stage of their natural history (Aronowitz 2007), by medical criteria and quantitative thresholds for diagnosis (Greene 2007), by the variable recognition that

certain diseases receive in medicine and medical statistics (Grob 2005), and by the decision about when to declare an epidemic or pandemic in the first place. But that only means that epidemics are a helpful mirror of medical priorities, concepts, and practices. What they reveal is that noncommunicable diseases—Omran's "degenerative and manmade diseases"—became a conceptual and material focus for modern medicine in the twentieth century.

Omran called the shift in morbidity and mortality from infectious diseases to noncommunicable diseases "the epidemiologic transition." As just mentioned, it was as much a transition in modern medicine's gaze toward new kinds of diseases and ways of combating them as it was a transition in morbidity and mortality. Even more radically, Omran's theory of the epidemiologic transition itself serves as an example of a related twentieth-century shift in modern medicine. "Epidemiology" names not just the subject of the epidemiologic transition—shifting patterns of disease—but also the science documenting the transition, one "concerned with the distribution of disease and death, and with their determinants and consequences in population groups" (Omran 2005, 731). Omran's is not only a theory of the epidemiology of disease but also an epidemiological theory (Krieger 2011). Without the growth and development of epidemiology in the twentieth century, the epidemiologic transition may have gone unanalyzed. Thus the classical epidemiologic transition hints at a wider transformation in twentieth-century medicine, one in which epidemiology itself infiltrated medical science as well as our understanding of disease and medical evidence. I call this three-part transformation in science, disease, and evidence a *threefold epidemiologic transition*.

Like "laboratory science," epidemiology is a science that encompasses many more specific sciences and disciplines, including infectious disease epidemiology, noncommunicable disease epidemiology, clinical epidemiology, and molecular epidemiology. These sciences each have their own specializations. What ties these fields together is partly their subject. Omran's description of epidemiology as studying the distribution and determinants of disease in human populations paraphrases a common way of characterizing the science and its subject (Krieger 2011). What it leaves out is that the sciences of epidemiology are equally characterized by their methods (Morabia 2004; Broadbent 2013). Epidemiologists *observe* exposures and outcomes in populations (and sometimes introduce exposures experimentally), *quantify* them at the population level using epidemiologic measures, and *compare* populations to make causal inferences about the determinants of

disease, exposures that affect human health such as medical interventions. In investigating disease, exposures, and the outcomes of intervention, epidemiology influences modern medicine's conception of all three, as we will see throughout this book.

While the recording of vital statistics and the making of population comparisons is a practice dating back to Britain in the 1600s and 1700s, respectively, epidemiologic methods and disciplinary identity matured in the twentieth century, especially after the Second World War (Susser 1985; Susser and Susser 1996; Morabia 2004).[23] The scope of epidemiology also expanded over the first half of the century with the investigation of not only infectious diseases but increasingly the new epidemics of noncommunicable diseases. Two methods characteristic of the maturing science of noncommunicable disease epidemiology in the 1900s were cohort studies and case-control studies. In a case-control study, exposures among cases afflicted with a health outcome are compared to exposures among unafflicted controls. In a cohort study, associations are studied in reverse: health outcomes among a population with a certain level of exposure are compared to health outcomes among a population with a different level of exposure. In both studies, an association found between exposure and a health outcome can be used to support the hypothesis that the exposure caused the health outcome.

Notable earlier twentieth-century examples include cohort studies of tuberculosis etiology (Frost 1939) and case-control studies of breast cancer etiology (Lane-Claypon 1926). After 1945 these methods and accompanying epidemiologic concepts were refined in the U.S. and Britain in research investigating the link between smoking and the rising rate of death among men from lung cancer. The British MRC supported two of the key studies: a case-control study among hospital inpatients (Doll and Hill 1950) and a prospective cohort study of British doctors begun in 1951 (Doll and Hill 1954), known as the British Doctors Study. Epidemiologic research provided the bulk of the data that led the U.K. Royal College of Physicians in 1962 and the U.S. surgeon general in 1964 to conclude emphatically that smoking caused lung cancer. Epidemiologists at the MRC also performed mathematical modeling that explained the age-related incidence of adult cancers

[23] The boundaries between epidemiology and "medical statistics," the application of statistics to health problems, are porous. Many contributors to the science of epidemiology were trained in statistics and were seen or saw themselves as contributing to medical statistics; they wrote articles or books on statistics and occupied positions as professional statisticians. Nonetheless, they simultaneously contributed to epidemiology when they worked to solve problems in public health using the characteristic methods of modern epidemiology.

and strengthened the hypothesized link between cancer and genetic mutation (Armitage and Doll 1954). Meanwhile, the British Doctors Study also supported a causal association between smoking and two other important noncommunicable conditions of the twentieth century, heart attack and chronic obstructive pulmonary disease (Doll and Hill 1956). The causes of heart disease were further enumerated in a cohort study of the community of Framingham, Massachusetts, the Framingham Heart Study (Kannel et al. 1961), begun in 1948. They came to include high blood pressure, high blood sugar, and high cholesterol.

Alongside new methods of noncommunicable disease epidemiology that interrogated the etiology of disease, clinical trials were refined in ways to better measure the effectiveness of interventions, including the new interventions of the postwar pharmaceutical revolution. Earlier clinical trials were often uncontrolled or matched trial groups based on observed characteristics, though in the first half of the 1900s the practice of allocating participants to trial groups in an alternating fashion was sometimes advocated (Marks 1997; Morabia 2004). The pivotal development for the twentieth-century history of epidemiology was the adoption of the randomized controlled trial (RCT) in medicine in the 1940s. In a standard RCT, trial participants are randomly allocated or randomized to two or more groups, which may be treated with new interventions, existing interventions, or a placebo or sham intervention. Health outcomes among trial groups are then compared to determine the comparative effect, if any, of the interventions trialed. The statistician R. A. Fisher (1935) had provided the statistical foundations of the randomized field experiment in the 1930s. The newly formed Therapeutic Trials Committee of the MRC then introduced randomization into its clinical trial design with RCTs testing two of the new postwar interventions: a vaccine for pertussis (MRC 1951) and the antibiotic streptomycin for pulmonary tuberculosis (MRC 1948).

The streptomycin trial finished first and is widely recognized as the first modern RCT and a model for future clinical trials (Matthews 1995). Soon after, the U.S. Public Health Service completed its own RCT of streptomycin for tuberculosis. Other publicly supported RCTs in America tested further interventions of the postwar revolution, including studies in the 1950s on cortisone for rheumatoid arthritis and rheumatic fever and on cancer chemotherapy for leukemia (the first randomized cancer clinical trial), as well as studies in the 1960s on antidiabetic agents and cardiovascular events and on antihypertensives for managing moderate hypertension

(Marks 1997; Greene 2007; Keating and Cambrosio 2007). Clinical trials testing antihypertensives and antidiabetic agents became the basis for a new pharmaceutical-based, patient-centered preventive medicine (Greene 2007; Aronowitz 2015). The wide scope of the public funding for studies of etiology and effectiveness in Britian and America illustrates the extent to which epidemiology became a state-supported science in the second half of the 1900s. Yet, while in the middle of the century public institutions were the largest sponsor of trials, by the end of the century industry outpaced the state in RCT publications (Bothwell et al. 2016).

The rise of RCTs and other clinical studies seeking to evaluate clinical practice was soon followed by the emergence of a new strain of epidemiologist specializing in them, the clinical epidemiologist (Daly 2005). Beginning in the late 1960s, the new science of clinical epidemiology advanced the methodology of RCTs and systematic reviews of the medical literature, eventually inventing standardized ways of designing and reporting this research. In the 1980s, it also offered up a new method for aggregating RCTs and other epidemiological studies, the meta-analysis, in which the results of multiple studies asking the same question are pooled to generate a combined effect estimate. The method allowed conclusions to be drawn about the effectiveness of entire classes of interventions, such as the new cholesterol-lowering statin drugs that were also introduced in the 1980s (Greene 2007).

Epidemiology thus developed as a response to two features of modern medicine: the rise in importance of noncommunicable diseases and a glut of medical interventions produced in the laboratory, both of which seemed to demand new methods of population analysis. In the process, epidemiology altered medicine's perspective on disease and evidence in ways that will be elaborated in subsequent chapters and mentioned briefly here.

In studying the diseases of the classic epidemiologic transition, noncommunicable disease epidemiology influenced medical ideas about diseases, their etiology, and their natural history. Alongside the laboratory, epidemiology provided data and modeling supporting the hypothesis that cancer is caused by multiple genetic mutations (instead of being caused universally by cancer viruses, for instance). Epidemiological modeling also explained why cancer incidence was increasing as populations aged, one aspect of the epidemiologic transition. By measuring multiple associations between noncommunicable diseases and exposures in case-control and cohort studies, epidemiology shook off a monocausal understanding of disease inherited from the germ theory of disease and animal

experiments, replacing it with a multifactorial understanding of disease in which (for instance) smoking causes both cancer and heart disease, while heart disease is variably caused by smoking and other factors. Finally, epidemiology facilitated the study of chronic diseases, those with a long duration and slow progression. While the methods of laboratory medicine were apt for studying acute diseases in animals through experimental manipulation, backward-looking (retrospective) and forward-looking (prospective) epidemiologic studies proved useful for studying the long latency, long progression, and long-term manifestations of chronic noncommunicable diseases.

Meanwhile, in evaluating medical interventions of the postwar revolution, clinical epidemiology influenced medical thinking on evidence and prediction as well as medicine's therapeutic attitude. Before reviewing how, it's necessary to introduce a movement that began near the end of the twentieth century, evidence-based medicine (EBM). Leading clinical epidemiologists who were convinced that the results of clinical epidemiology evidence provided the proper evidence base for clinical practice formed the EBM Working Group, centered at McMaster University in Canada (Daly 2005). In the early 1990s, members of the group described EBM as "[a] new paradigm for medical practice" that elevated evidence from clinical epidemiology above clinical experience and pathophysiologic reasoning for grounding medical decision-making (EBM Working Group 1992, 2420).[24] For therapy, the RCT and systematic review of RCTs (often including a meta-analysis) was the pinnacle of high-quality evidence. EBM created "rules of evidence" for medicine—most notably, "hierarchies of evidence," or evidence grading schemes (e.g., Balshem et al. 2011)—that operationalized its stance on the quality of different kinds of evidence. The principles of clinical epidemiology and EBM's rules of evidence were quickly incorporated into medical training, journal publishing policies, standards for systematic review and clinical guideline development, and, eventually, institutions beyond medicine. Perhaps the most revealing definition of EBM and its epidemiological roots is provided by a popular EBM

[24] "A new paradigm for medical practice is emerging. Evidence-based medicine de-emphasizes intuition, unsystematic clinical experience, and pathophysiologic rationale as sufficient grounds for clinical decision making and stresses the examination of evidence from clinical research. Evidence-based medicine requires new skills of the physician, including efficient literature searching and the application of formal rules of evidence evaluating the clinical literature" (EBM Working Group 1992, 2420).

textbook: "EBM is the use of mathematical estimates of the risk of benefit and harm, derived from high-quality research on population samples, to inform clinical decision-making in the diagnosis, investigation or management of individual patients" (Greenhalgh 2010, 1).

Clinical epidemiology thus shaped medicine's understanding of what counts as medical evidence, to the point that "evidence" in medicine became virtually synonymous with the results of clinical epidemiology studies. Rather than making therapeutic predictions from unsystematic patient trials (clinical experience), from the results of laboratory studies, or from pathophysiologic theory, clinical epidemiology enjoined clinicians to extrapolate from the results of RCTs and other clinical epidemiology studies. Finally, clinical epidemiology altered scientific medicine's therapeutic attitude from one of optimism to one of skepticism. While laboratory science produced a rash of drugs, clinical epidemiology tested them to determine their effectiveness, relying on its methods to minimize and evaluate bias lurking in the evidence on those interventions. Skepticism seemed necessary to temper the inclination of clinicians to intervene and the inclination of drug companies to profit. As we will see momentarily, this effort included regulatory schemes that raised RCTs as a barrier to the medical marketplace.

On the border between disease and evidence, epidemiology also altered the object of medical diagnosis, prognosis, and intervention with a newfound focus on risk. The etiologic factors it identified in observational studies became known as "risk factors" because they foretold a quantifiable risk of disease. Quantifying this risk became a task for medical prognosis. Asymptomatic risk factors such as borderline high blood pressure, raised blood sugar, and moderately high/abnormal cholesterol themselves were henceforth seen as diagnosable diseases (hypertension, non-insulin-dependent diabetes, hypercholesterolemia/dyslipidemia). Managing and preventing incurable chronic conditions took on increased significance. And "risk reduction," measured in RCTs, became a common goal of prevention and therapy, whether in controlling disease risk factors or combating cancer.

The threefold epidemiologic transition in disease, evidence, and science so changed modern scientific medicine that it is no longer fitting today to speak of laboratory medicine as the medical model. The story of modern medicine, told as a succession of medical models, needs updating.

## 1.6. The New Medical Model: Epidemiological Medicine

The age of modern scientific medicine started with the rise of laboratory medicine, the medical model. Yet in the postwar period, medicine's threefold epidemiologic transition gave rise to a new modern medicine, one best characterized not as laboratory medicine nor as biomedicine but as epidemiological medicine.

Bernard and Virchow are often associated with nineteenth-century scientific medicine, as figures who envisioned a laboratory science–based medicine, one integrated with sciences like physiology and pathophysiology. They might also be seen as intellectual founders of the modern medical tradition of today, characterized by the intertwining of clinical medicine and medical science. Yet to understand contemporary scientific medicine and the ways it has evolved since the nineteenth century and especially since the Second World War, their vision and contributions are less salient than that of the epidemiologist Austin Bradford Hill, who advanced the design of clinical trials, the methods and principles of noncommunicable disease epidemiology, and the cause of epidemiology in medicine, all within a few short years (Doll 1992).

Hill had no formal training in medicine, epidemiology, or statistics. Nonetheless, as director of the MRC's Statistical Research Unit, he proposed the use of randomization in its trials of streptomycin for tuberculosis and vaccination for pertussis and oversaw its use in subsequent trials of cortisone for rheumatoid arthritis and rheumatic fever. For Hill, randomization helped the trialist overcome bias, that central concept of modern epidemiology. Hill also led the MRC's case-control and cohort studies of smoking and lung cancer, the latter of which suggested that smoking was a cause of further noncommunicable diseases. Hill's (1965) principles for inferring causation from association in observational epidemiological studies, first applied in the context of this research, became canons for the developing science of epidemiology. Finally, Hill brought epidemiological thinking to a medical audience through a series of articles in the medical journal *The Lancet* in 1937, published that same year in his authoritative textbook *Principles of Medical Statistics*, a guide to epidemiological thinking in medicine (Morabia 2004). Hill's *Principles* serve as a record of the development of epidemiology and its changing place in medicine during the postwar years. Over eleven subsequent editions, Hill incorporated his ideas on inference and methodology

in clinical trials and observational epidemiology. Through his research and advocacy for epidemiology in medicine (Matthews 1995), Hill helped steer medicine through its threefold epidemiologic transition.

In a pair of articles in the *New England Journal of Medicine* on the clinical trial (1952) and on observational epidemiology (1953), Hill argued that the epidemiologist's "statistical approach" was useful for "the practiced arts of the clinician" (1952, 113) and not only for public health, and that controlled observations in humans were the ultimate way to settle questions of causation when it came to disease etiology and treatment effectiveness; they provided "the final proof" (1953, 999) above laboratory experiments on animals and uncontrolled observations in humans.[25] Hill (1953, 1000) quoted the Oxford physician George Pickering: "Any work which seeks to elucidate the cause of disease, the mechanism of disease, the cure of disease, or the prevention of disease, must begin and end with observations on man."

Hill's model of scientific medicine was not laboratory medicine; it was one in which clinical medicine was integrated with and modeled after epidemiological research, which took on greater authority as a source of medical knowledge than the results of laboratory experimentation or uncontrolled trials in patients. Adding to the tradition that names models of scientific medicine after their characteristic sciences or scientific locations—Paris *clinical* medicine (or *hospital* medicine), *laboratory* medicine (Ackerknecht 1967; Jewson 1976)—I call the new scientific medicine *epidemiological medicine*.[26]

Epidemiological medicine is integrated with the sciences of epidemiology. First, it integrates the results of clinical epidemiology evidence on diagnostic tests, prognostic markers, and medical interventions into patient care. In fact, according to the most common definition of EBM, "The practice of evidence based medicine means *integrating* individual clinical expertise with the best available external clinical evidence from systematic research" (Sackett et al. 1996, 71, my emphasis). EBM determines what counts as "the

[25] Careful not to alienate his medical audience and denigrate other methods of investigation beyond trials, Hill also defended the importance of the clinical impression (1952) and nonexperimental epidemiological observations (1953). In this vein, he ended is article on the clinical trial on a softer note: "The statistically guided therapeutic trial is not the only means of investigation and experiment, nor indeed is it invariably the best way of advancing knowledge of therapeutics. I commend it to you as *one* way, and I believe, a useful way" (1952, 119).

[26] Alternatives such as "evidence-based medicine," "risky medicine" (Aronowitz 2015), and "desktop medicine" (Karlawish 2010) capture only some of the ways that the new scientific medicine is integrated with and modeled after epidemiology. In what Jason Karlawish (2010) calls "desktop medicine," a networked computer on a desk replaces a patient in a hospital bed as the center of medical attention. The computer is epidemiological medicine's portal to the medical literature, including the epidemiological literature.

best available external clinical evidence," namely, the results of clinical epidemiology studies. EBM proponents saw clinical epidemiology as the science of clinical care, and "scientific medicine" was reportedly even one of the initial names proposed for their movement (Daly 2005). Epidemiological medicine also integrates knowledge about the etiology and natural history of diseases from observational epidemiology, including noncommunicable disease epidemiology. Epidemiology, as a science that studies the distribution and determinants of health and disease in *populations*, is most naturally and commonly viewed as a public health science. In the second half of the 1900s it also became an important science for clinical medicine, which required medicalizing its research and results and making them palatable to patient care, using, for instance, individualized concepts of etiology and risk (Pearce 1996; Krieger 2011; Phillips 2025).

As with earlier scientific medicines, some of the integration in epidemiological medicine is also institutional, as hospitals became convenient sites for epidemiological research, especially clinical trials. Early on in the discipline, many epidemiologists were physicians, and some still are. Hill also advocated for research collaboration among epidemiologists and clinicians, which soon come to pass (Matthews 1995; Marks 1997). Moreover, clinicians became deeply aligned with public health epidemiology as an arm of the public health workforce, administering childhood vaccinations and advising adults to stop smoking and avoid other risk factors for heart disease and cancer. Later EBM styled itself as a new paradigm for the institution of medical education, planting itself in the soil from which new generations of doctors would rise. EBM principles now influence standards of medical care, including (evidence-based) clinical guidelines and medical insurance coverage, as care must be based on the best available evidence (from clinical epidemiology). Consensus clinical guidelines produced by medical organizations and setting standards of medical care became more common and influential in the latter part of the 1900s (Timmermans and Berg 2003; Weisz et al. 2007; Upshur 2014; Solomon 2015), emerging as a major vehicle for integrating medicine with epidemiology and infusing medicine with epidemiological thinking.

With epidemiological medicine, scientific medicine also underwent "remodeling," as clinical medicine became modeled after epidemiological sciences that transformed its perspective on disease and evidence. Diseases are now understood in terms of risks and molecules. Their etiology is multifactorial, consisting of the multiple risk factors, including

genetic causes, measured in epidemiological studies. Their pathology is more often nonspecific (rather than necessarily defined by specific germs or physiological disturbances). Diseases are frequently chronic and asymptomatic, with a long progression studied by following population cohorts. The goal of clinical care: to prevent, treat, and manage diseases and their manifestations through interventions shown to reduce the risk of outcomes in epidemiologic studies,[27] particularly RCTs, which serve as medicine's evidence base. Therapeutic inferences involve extrapolating from the results of these epidemiological studies. And the therapeutic attitude is primarily skeptical, as the medical evidence is scrutinized for signs of bias.

Table 1.1 compares epidemiological medicine to laboratory medicine, displaying their main differences. A crucial difference between laboratory and epidemiological medicine is that while laboratory medicine casts clinical statistics ("the numerical method") as an inferior source of medical knowledge compared to knowledge gained from the laboratory, epidemiological medicine takes the opposite view, privileging clinical and community research.

Some have drawn comparisons between EBM or clinical epidemiology on the one hand, and Paris clinical medicine or the numerical method on the other (Sacket et al. 1996; Rangachari 1997; Weisz 2005; Morabia 2006), which is fitting because epidemiological medicine shares with Paris clinical medicine an emphasis on population comparisons and quantification. There are also important differences. Of course, epidemiological medicine favors methods such as RCTs that were not in use during the nineteenth-century heyday of Paris clinical medicine. (Louis mainly compared hospital case records to determine the comparative effectiveness of interventions.) An emphasis on pathological anatomy and the physical exam in Paris

[27] In fact, in epidemiological medicine, an intervention can be defined as a preventive or curative medical intervention when it reduces the risk of a medical condition in a population (Chapter 6). Cure remains a central goal of medicine (Broadbent 2019); the point is that in epidemiological medicine, its achievement is often documented through risk reduction in epidemiologic studies. Simultaneously, it is accepted that cure is often out of reach for incurable chronic conditions, elevating the goals of disease prevention and management in a new pharmaceutical- and individual-oriented preventive medicine. Hence prevention, management, and other forms of treatment—conceived as risk reduction—are the noteworthy goals in "risky" epidemiological medicine (Aronowitz 2015). I will set aside the question of whether medicine's concepts of health and disease were also remodeled in epidemiological medicine from ones defined by physiological function and dysfunction (Boorse 1977) to ones defined by epidemiological concepts such as risk (Giroux 2015), though I examine examples of "risk-based diseases" (Schwartz 2008) in Chapter 4 and I examine risk in Chapter 6. Regardless of whether risks and risk factors redefine "disease" in epidemiological medicine, they are undoubtedly central objects of diagnostic, prognostic, and therapeutic concern.

**Table 1.1** Comparisons Between the Laboratory Medicine and Epidemiological Medicine Models

| | **Laboratory Medicine** | **Epidemiological Medicine** |
|---|---|---|
| Research Sciences | Experimental laboratory sciences (e.g., physiology, biochemistry, bacteriology) | Epidemiological sciences (e.g., noncommunicable disease epidemiology, clinical epidemiology) |
| Explanation | Pathophysiological | Molecular |
| Etiology and Pathology | Specific etiology/pathology | Multifactorial etiology, nonspecific pathology |
| Clinical Course | Acute, often symptomatic | Acute or chronic, often asymptomatic |
| Diagnosis | Precision instruments, pathophysiology | Risk factors, risk |
| Therapeutic Goals | Eliminate pathogen, restore normal function | Prevent/manage disease, reduce risk |
| Therapeutic Evidence | Experimental (laboratory, bedside) | Epidemiological (e.g., RCTs) |
| Therapeutic Reasoning | Rational therapeutics, mechanistic reasoning | Extrapolation from epidemiologic studies |
| Therapeutic Attitude | Optimistic | Skeptical |

clinical medicine has recently been replaced by an emphasis on molecular diagnostics. And concepts of risk are central to epidemiological medicine but were not used by followers of Louis, who often invoked the "calculus of probabilities" but rarely used probabilities or anything more than rudimentary statistics (Matthews 1995). Nonetheless, there is a certain kinship between Paris clinical medicine and EBM. In Chapter 5, I argue that they are best seen as two varieties of empiric medicine because of their reliance on inferences from clinical observations rather than pathophysiological theorizing.

In ways that I will articulate momentarily, epidemiological medicine has replaced laboratory medicine, the original medical model, as the favored model for orthodox scientific medicine. As orthodox scientific medicine remains powerful—if not more powerful—in society, epidemiological medicine is rightly considered *the new medical model*, the newly dominant model of medicine in society. If modern scientific medicine began with the rise of laboratory medicine, then the new medical model of epidemiological

medicine marks a *new modern medicine*, one shaped by new disease epidemics, new norms of evidence, and the new sciences of epidemiology.

John Burnham (2015) advances a different history of modern medicine in which scientific developments in each new medical era are layered on top of preceding developments, beginning with a "modernizing epoch" in the 1880s initiated by the laboratory.[28] There is something that this complementary story gets right. The laboratory medicine model remains relevant in contemporary medicine. The period from the 1940s to the 1960s was one of extensive overlap in the influence of both laboratory and epidemiological models, as a time that witnessed both the pharmaceutical revolution and the birth of epidemiology as we now know it. (Following the 1960s, the pace of therapeutic innovation seemed to slow; Le Fanu 2012). Physiological concepts, knowledge, and reasoning are still used in practice today, alongside laboratory diagnosis and laboratory therapeutic monitoring. The clinical deployment of drugs, biologicals, and biochemical assays has only increased and is often still justified directly through the results of laboratory studies or simple bedside experimentation. And laboratory medicine may have special purchase in the treatment of rare or novel disorders for which there is little population data to guide clinicians, and all they are left with are physiological first principles and trial and error. Meanwhile, laboratory science continues to flourish as a research tradition, contributing massively to the growth of scientific knowledge.

To some extent, epidemiology is an addition to rather than a substitution for the scientific palate of contemporary scientific medicine, and in arguing that scientific medicine today is integrated with and modeled after epidemiology I am emphasizing what is novel about contemporary scientific medicine. However, there is something important that Burnham's historical layering leaves out, something captured by my picture of models waxing and waning in influence. As laboratory medicine eclipsed Paris clinical medicine in the 1800s, epidemiological medicine eclipsed laboratory medicine in the 1900s in clinical priority. Recall that by "clinical priority," I mean the comparative power and influence that particular models and their respective sciences have on clinical practice. Because scientific medicine is clinical medicine integrated with and modeled after science, clinical priority is the central determinant of which science best characterizes the dominant model

[28] Burnham (2015) also argues that a new medical epoch began in the 1980s and is best characterized by a transformed clinical encounter in which the "sick role" is replaced by the "at-risk role" (Burnham 2012) and doctors practice "desktop medicine" (Karlawish 2010) rather than bedside medicine.

of scientific medicine at a given time. In the second half of the 1900s, epidemiology began to enjoy greater clinical priority than laboratory science, particularly when it came to clinical reasoning and decision-making, and especially in justifying the clinical use of treatments.

Before the Second World War, a less regimented mix of clinical and laboratory studies was often used to support the new treatments. In the postwar period, through various institutions and developments, there emerged an order that accorded greater justificatory weight to statistical clinical studies, especially RCTs (Marks 1997), placing them in greater proximity to the clinic, both temporally and epistemologically (Figure 1.1). We already encountered one of these developments, the EBM movement of the 1990s and its hierarchies of evidence. The hierarchies are schemes that grade evidence according to its quality. They are often presented in the form of an evidence pyramid, with higher quality evidence found above lower quality evidence (Figure 1.1A). The highest quality evidence is the best evidence for clinical decision-making. Though there have been many EBM hierarchies offered by different groups, they have in common a ranking of epidemiological evidence above non-epidemiological evidence, including evidence from the laboratory, clinical experience, and pathophysiological/mechanistic reasoning (Howick 2011).

A second phenomenon that cemented a greater clinical proximity for epidemiological research compared to laboratory research was the creation of a system of phased research for drug development, which divides research into preclinical and clinical phases. Preclinical research includes drug discovery and development as well as in vitro and animal testing in the laboratory. Clinical research includes a graduated series of phase I, phase II, and phase III clinical trials involving successively more participants. Phase III trials are typically randomized. Thus we find laboratory science at one end and epidemiology at the other. Preclinical research typically takes place before clinical trials and directly informs clinical trials, while phase III trials directly inform clinical practice. Protocols for phased drug development were used by the U.S. National Cancer Institute and the European Organization for Research and Treatment of Cancer in the 1960s (Keating and Cambrosio 2007). The Food and Drug Administration (FDA) took the lead worldwide in establishing the phase III randomized drug trial as the gold standard for market approval between 1962 and 1970 (Bothwell et al. 2016), adopting the phased approach as a regulatory standard (Figure 1.1B).[29]

[29] On the FDA's power to influence medical concepts and epistemology, see Carpenter (2010).

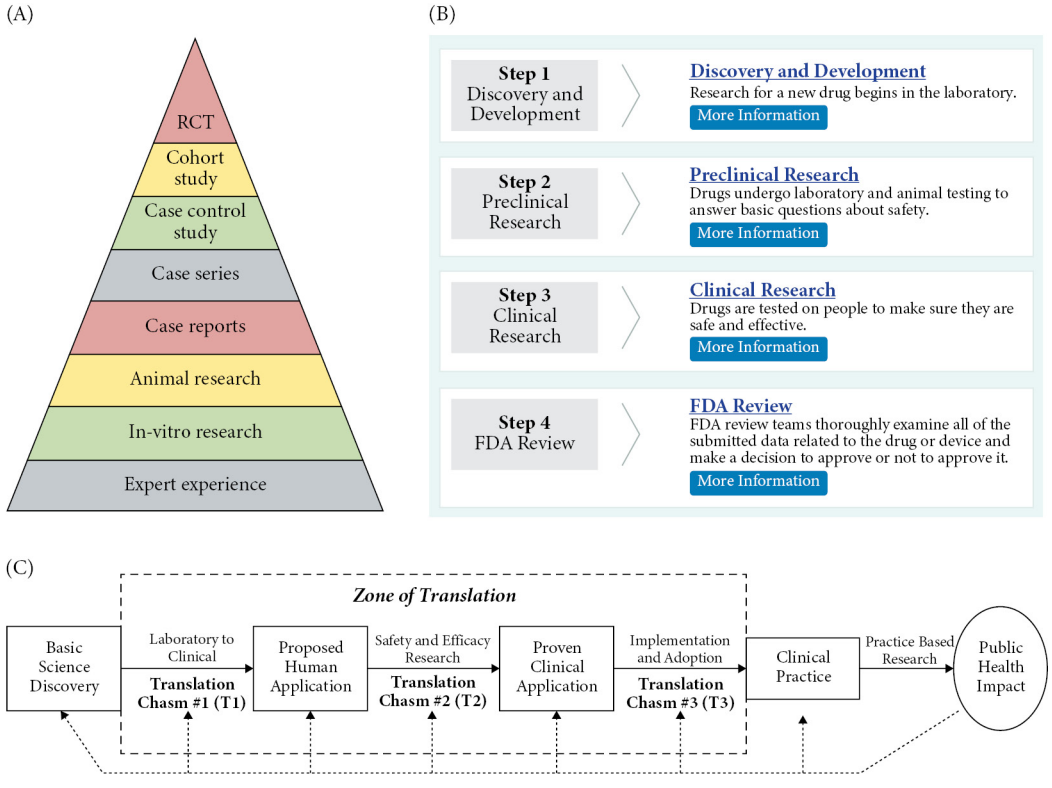



**Figure 1.1** The greater proximity of epidemiological science to practice compared with laboratory science in contemporary scientific medicine. A: the traditional EBM hierarchy of evidence; B: the drug development process; C: the translation continuum.

A: Djulbegovic Guyatt (2017, 417); B: Food and Drug Administration (2018); C: Drolet and Lorenzi (2011, 4).

A third and related two-part development was the growing distinction between basic and applied science and the later concept of the research translation continuum. In the postwar period, the NIH and other federal agencies distinguished "basic" from "applied" science, where basic science was understood to be research into "fundamental biological processes" (NIH Study Committee 1965, 2), or mechanisms of health and disease. On the other hand, in applied clinical research, clinical applications resulting from basic discoveries are studied and evaluated in clinical trials, including RCTs, the results of which can then directly inform health policy and clinical decision-making (Sung et al. 2003; Butler 2008). While clinical trials and other epidemiological studies might contribute to basic science,[30] laboratory research is *not* included in the category of applied clinical science and thus is, conceptually speaking, further from the clinic.

By the 2000s, there arose concern at the NIH and in the broader medical research community that research was not yielding effective interventions as efficiently as expected. Basic research had to be better "translated" into successful discoveries (Zerhouni 2003; Butler 2008). The "translational medicine" or "translational research" movement arose to try to improve this process (Solomon 2015). The "translation continuum" describes various gaps in the research-to-practice or "bench-to-bedside" process to be bridged by research and initiatives, including a gap between basic research and interventions that make it all the way through clinical trials, and a subsequent gap between trial results and changes in clinical practice (Figure 1.1C). The metaphor of translation is revealing of our contemporary model of scientific medicine. In contemporary scientific medicine, the results of laboratory research must first be "translated" into a proven intervention through epidemiological "safety and efficacy research" (clinical trials). The results of this research, written as they are in the language of the clinic, can then be implemented in clinical practice, along with the intervention.[31] In laboratory medicine, there was no need to translate experimental laboratory results because they revealed universal laws of physiology and pathophysiology

[30] The extent to which epidemiological studies contribute to a basic understanding of mechanisms is controversial. Studies in chronic disease epidemiology and clinical epidemiology are often said to follow a model of "black box" research by bracketing the mechanisms between exposure/intervention and outcome (Susser and Susser 1996; Broadbent 2011). Yet some insist that clinical research can, does, or should contribute to basic knowledge (Keating and Cambrosio 2007; Hey 2015; Solomon 2015; Keating et al. 2016).

[31] In Figure 1.1C, "practice based research" includes further epidemiologic studies: research using patient registry data, cohort studies, case-control studies, and meta-analysis (Drolet and Lorenzi 2011).

that could be exploited by therapeutic principles to provide direct rational grounds for therapy. In sum, since the Second World War various institutions have sought to order research methods using conceptual models that place epidemiology in closer proximity to clinical practice compared with laboratory science.

In contrast, Miriam Solomon (2015) describes contemporary medical methods, including the methods of EBM and translational research, as constituting a "developing, untidy, methodological pluralism." It is an *untidy* pluralism, according to Solomon, because "sometimes the methods are applied to answer the same questions [such as questions about the effectiveness of a medical intervention], giving sometimes the same and sometimes different answers, and there is no 'meta-method' to which to appeal to resolve the disagreements" (208), "no methodological shortcuts (such as general methodological hierarchies) to resolving controversy" (224). In contrast, "[i]f the methods *always* worked in different domains or for different questions, we would have . . . a 'tidy pluralism'; each method could have its own domain of application coinciding with its domain of authority" (208, my emphasis). As an example of relatively unmanaged disagreement among medical experts, Solomon discusses the controversy over the effectiveness of screening mammography for women in their forties.

But, *pace* Solomon, the EBM hierarchy of evidence *is* a widely accepted "meta-method" or "methodological hierarchy" for deciding which kind of evidence to believe—whether evidence from basic laboratory science or clinical epidemiology evidence—in answering questions about the effectiveness of medical interventions, while the phased drug development approach and the translation continuum assume a kind of ordering of basic laboratory research and applied clinical research. There can always be exceptions that defy these frameworks.[32] Nonetheless, the recent history of medicine is one of institutional efforts to order medical research methods into hierarchies, phases, and continua. Thus I conclude that modern medicine's developing scientific pluralism is relatively tidy, even if epidemiological science does not *always* have clinical priority over laboratory science.

An important alternative to my view—that epidemiological medicine is the new medical model of the postwar era—to consider is perhaps the received view: that the postwar medical model is biomedicine or the

[32] Solomon (2015) notes that the screening mammography case is "atypical" and that "evidence-based approaches" are commonly privileged.

biomedical model. In 1977, in an article in *Science* titled "The Need for a New Medical Model: A Challenge for Biomedicine," the psychiatrist George Engel wrote, "The dominant model of disease today is biomedical" (130). Engel suggested that biomedicine was a "scientific model," "devised by medical scientists" (382). According to Engel, the biomedical model embraced "reductionism," understanding disease fully through biological variables and excluding the psychological and social dimensions of illness. This predisposed medicine toward interventions targeting biological variables. The biomedical model, medical model, or biomedicine became a locus for diverse criticisms of medicine, each having more or less to do with Engel's characterization but often scrutinizing mainstream medicine's tendency to prioritize the biological, to the neglect of the psychological and social.[33] Peter Keating and Alberto Cambrosio (2003, 2004) take biomedicine to describe a greater "alignment" between biology and medicine in the postwar era.[34]

If we treat biomedicine as a model of scientific medicine akin to laboratory medicine and Paris clinical medicine, then what kind of research is the "biomedical science" that clinical medicine is integrated with and modeled after in the biomedical model? For his part, Engel claimed that molecular biology was the biomedical model's "basic scientific discipline" (1977, 379). However, the meanings of "biomedical science" or "biomedical research" are varied.

Though one can find isolated earlier mentions of the terms "biomedical science," "biomedical research," and "biomedicine," the terms and concepts took hold in medical research only after the Second World War. They are most associated with the NIH, particularly due to a landmark report on the NIH published in 1965 by Dean Wooldridge titled "Biomedical Science and Its Administration" (NIH Study Committee 1965), commonly known as the Wooldridge Report (Keating and Cambrosio 2004; Roth 2022). The Wooldridge Report observed that much research across the

[33] Engel (1977) believed that psychiatry was facing a crisis because it did not conform to the medical model. However, rather than part ways, Engel argued that both medicine and psychiatry should embrace a biopsychosocial model of disease. The biopsychosocial model has also become the subject of an enormous body of scholarship in the humanities and in medicine. For philosophical perspectives on the biomedical and biopsychosocial models in medicine and psychiatry, see Murphy (2011) and Gifford (2017).

[34] Other features associated with "biomedicine" after the Second World War by historians and sociologists include greater collaboration among biologists, clinicians, and industry; a closer connection between the biological laboratory and the clinic; a rapid expansion of funding for biomedical research; increasing application in the clinic of therapies developed in the laboratory; and the "molecularization" of medical science and medicine (Gaudillière; 2002; Keating and Cambrosio 2003, 2004; Quirke and Gaudillière 2008; Löwy 2011).

NIH's intramural and extramural divisions was basic, studying biological mechanisms, especially through experimental laboratory research. A great deal of it was devoted to the new science of molecular biology that Engel credited as biomedicine's basic science. Yet as the Report notes, the range of basic research supported by the NIH was in fact much wider (and included the older sciences of laboratory medicine).

When it came to stating plainly the purpose of the NIH, the Report enlarged the scope of its research even wider, concluding that "the primary de facto mission of NIH is the stimulation and support of a very broad range of health related or biomedical research" (NIH Study Committee 1965, 3). While "biomedical science" sometimes refers more narrowly to basic laboratory research (National Research Council 2011b), its meaning has loosened over time to encompass the "broad range of health related research" described in the Wooldridge Report (Keating and Cambrosio 2007; Roth 2022), including epidemiological research.[35]

Landmark epidemiological studies took place under the banner of the NIH, including many of the clinical trials mentioned in the previous section: trials of streptomycin, cortisone, antidiabetic agents, and cancer chemotherapy. The NIH's National Cancer Institute organized the first randomized cancer trial (of chemotherapy for leukemia) and became the largest sponsor of cancer trials in the U.S. The NIH also supported the Framingham Heart Study on the epidemiology of heart disease. Finally, NIH statisticians and epidemiologists refined epidemiological methods, setting the stage for the science of clinical epidemiology (Phillips 2020). The statistician Jerome Cornfield was an American Bradford Hill who spent much of his prodigious career at the NIH. Cornfield contributed to the analysis of case-control studies (Cornfield 1951), to multivariate risk analysis in Framingham (Truett et al. 1967), to the analysis of confounding in epidemiologic studies (Cornfield et al. 1959), and in various ways to clinical trials (Ederer 1982). He also published an influential synthesis of the epidemiological evidence on smoking and lung cancer (Cornfield et al. 1959). Meanwhile the NIH's Thomas Chalmers advanced the methods of RCTs and meta-analysis (Dickerson and Chalmers 2015).

[35] Nancy Krieger (2011) argues that epidemiology began to follow the biomedical model as it searched for multicausal biological and behavioral risk factors for disease. Keating and Cambrosio (2007, 222) refer to cancer clinical trials as a new style of practice incorporating elements of distinct laboratory and statistical styles and that "typifies modern biomedicine."

Therefore, in order of increasing breadth, the "biomedical science" of biomedicine could refer to molecular biology, basic laboratory science, or any health-related scientific research (including epidemiology). However, none of the three resulting models of scientific medicine adequately represents the new medical model of the postwar era. While molecular biology flourished after the war, its impact on clinical practice was sluggish and mostly limited to genetic diagnosis, at least initially (Burnham 2015). Moreover, epidemiology also advanced medicine's understanding of the genetic basis of disease (Chapter 2) and is emerging as the science of gene-centric precision medicine in the twenty-first century (Chapter 9); so sciences beyond molecular biology and including epidemiology have contributed to the "molecularization" of clinical medicine associated with the biomedical model.

The second possibility, that biomedicine is basic laboratory science–based medicine, is a reinvention of nineteenth-century laboratory medicine (Löwy 2011). While acknowledging this continuity, Viviane Quirke and Jean-Paul Gaudillière suggest that "biomedicine was only truly 'invented' after the Second World War" (2008, 442), by which they mean that it was only then effectively realized with a step change in the scale of investment in laboratory science and in its impact on patient health. While the scale is indeed remarkable, here I am interested in the emergence and rise in influence of new models for clinical practice. In this sense, the new invention of the postwar years was *not* laboratory medicine.

The third possibility, that biomedicine is clinical medicine integrated with any health-related scientific research, is either a different name for what I've been calling "scientific medicine" or a model of scientific medicine that, unlike previous models, does not prioritize particular research sciences. If the former, then biomedicine is not a new medical model unique to the recent past. If the latter, then biomedicine fails to account for the structured ordering of medical sciences in the recent past.

All three possibilities for biomedicine as a model of scientific medicine underemphasize the rise of epidemiology and its clinical priority in the postwar era. "Biomedical science" is best seen for what it originally was: a way of administrating publicly funded biological research that covered research with perceived potential health benefits and that excluded research without them (Roth 2022). While biomedical science is a major *research* priority in the new modern medicine, its research results have lesser *clinical* priority compared to epidemiology. Meanwhile, "the biomedical model"

best describes an emphasis on biological variables in *multiple* iterations of scientific medicine and the tendency to "reduce" the patient to their parts (with successive models tending to reduce them to smaller and smaller parts—most recently, to molecules). Early twentieth-century scientific medicine was just as guilty of prioritizing biological factors—in its case, bacteria, hormones, and other objects of physiology and biochemistry. The importance of these forces notwithstanding, a model of scientific medicine appeared in the second half of the twentieth century that deserves attention equal to that heaped on biomedicine. That model is epidemiological medicine, and it is the new medical model.

## 1.7. Conclusion

Modern medicine. Laboratory medicine. Biomedicine. These names for contemporary orthodox medicine are commonplace. Over the past two hundred years, medicine discovered a trove of effective medical interventions, culminating in a pharmaceutical revolution in the decades after the Second World War. Yet what made medicine modern over this two-hundred-year period was another revolution, a scientific revolution that saw the rise of modern scientific medicine. In modern scientific medicine, clinical medicine is integrated with and modeled after modern research sciences. One of these models of scientific medicine, a laboratory medicine modeled after experimental laboratory sciences, became *the* medical model by the early twentieth century, achieving dominance in society. Then, over the twentieth century, medicine underwent a threefold epidemiologic transition in disease, evidence, and science, in which all three were transformed by the new sciences of epidemiology. While the received view is that the biomedical model or biomedicine came to dominate medicine in the second half of the twentieth century, this narrative obscures the clinical priority achieved by epidemiological sciences and by a model of scientific medicine embracing them, epidemiological medicine. Epidemiological medicine is the new medical model. The threefold epidemiologic transition and the new medical model mark a new modern medicine that will be explored in the chapters to come.

# PART 2
# DISEASE

# 2

# Theories of Contagion and Cancer


## Summary

Are there scientific theories of disease in medicine? Some historians and philosophers seem to think so. What they often have in mind are not grand "theories of everything" explaining all disease but rather more modest theories about groups of diseases. What do these theories do? And what qualifies them as theories *of disease*? To answer these questions, I use as my prototype for a theory of disease the germ theory, an innovation of early modern scientific medicine. Considering that cancers are diseases characteristic of the *new* modern medicine, I then investigate whether there are theories of cancer in germ theory's image. Germ theory and somatic mutation theory (SMT)—as theories of contagion and cancer, respectively—unify models of the pathogenesis of diseases within a family of disease.

When laboratory scientists sought to extend the germ theory to cancers, this research resulted in the development in the late 1900s of a somatic mutation theory of cancer that saw cancer as a disease of uncontrolled cell proliferation, genetic alteration, and molecular dysregulation. While accurate, this story overlooks the role of epidemiological observations and an epidemiological theory, the multistage theory, in providing earlier support for the genetic origins of cancer. The multistage theory explained the increasing incidence of adult cancers with age by modeling carcinogenesis as a multistage process. Rather than theories of disease, some historians and philosophers have emphasized the importance of frameworks that organize and direct infectious disease research and cancer research. But theories of disease can play a similar role by providing an etiological blueprint, a causal formula, for studying disease. Last, the centrality of molecular explanations in contemporary medicine is epitomized by SMT, and this makes the theory susceptible to the charge of molecular reductionism: reductive explanations that explain disease through the interaction of genes and molecules can fail to capture the complexity of cancer and can exclude important social and environmental



*The New Modern Medicine*. Jonathan Fuller, Oxford University Press. © Oxford University Press 2025.
DOI: 10.1093/9780190066178.003.0003

causes. While the first half of this criticism is misdirected at SMT, the second half is closer to the mark.

## 2.1. Contagion and Cancer

Although earlier in the 1800s pathological anatomy and physiology enjoyed epistemological primacy in medicine, by the late 1800s bacteriology was the height of the medical sciences. Infectious diseases—both endemic diseases that claimed lives with cruel consistency and epidemic diseases that claimed them with frightening inconsistency—were among the deadliest and most feared maladies. In the second half of the century, laboratory medical science discovered the causes of several contagions, including tuberculosis and cholera, and provided hope for their control. The germ theory of disease, which attributed the cause of contagious diseases to microscopic germs, summarized these much-heralded achievements in the history of medicine.

Frank Snowden (2020, 204–205) holds back no embellishments when he writes of medicine at the end of the 1800s, "A scientific discipline whose basic principles were recognizably similar to those prevailing today had begun to emerge. . . . The closing decades of the nineteenth century—roughly between 1860 and 1900—witnessed a wholesale revolution in medical philosophy with the germ theory of disease as its central feature. With only slight hyperbole, one could argue that this theory constituted as important a revolution in medicine as the Copernican theory of the rotation of the earth around the sun was to astronomy, as gravitational theory was to physics, or as Darwin's theory of natural selection was to biology." Not only is the early history of germ theory enmeshed with therapeutic discoveries and developments such as rabies vaccine, diphtheria antitoxin, the anti-syphilitic agent Salvarsan, as well as changes in both personal and public hygiene, but, Snowden argues, it also helped establish the laboratory rather than the hospital ward as the center of medical authority, spurred reforms in scientific medical education (including an emphasis on experimental laboratory sciences), and helped scientific medicine and the laboratory medicine model secure their cultural authority. In a similar spirit, David Wootton (2007, 227) argues that Joseph Lister's application of germ theory to the problem of surgical wound infections in 1865, resulting in the discovery of antiseptic agents, marks the start of modern medicine because germ theory was a scientific theory: "Lister thus begins the modern history of medicine, defined in terms

of constant improvements in therapy grounded in constantly developing scientific understanding."

If germ theory deserves even a fraction of this esteem, it suggests—perhaps surprisingly—that theories of disease can play an important part in the evolution of modern medicine. I say "surprisingly" because theories of disease might seem to be a decidedly premodern fantasy. The idea that a grand, sweeping theory of disease could account for the motley of maladies that afflict us is probably most associated with Hippocratic medicine (Nutton 2012). The centerpiece of Hippocratic medicine was humorism, or the theory of the four humors, in which the balance or imbalance of four bodily humors determined health or disease, respectively. Hippocratic humorism was exactly the kind of "metaphysical" system that modern scientific medicine saw itself as upending (Warner 1998). However, germ theory can serve as a prototype for a modern *scientific* theory of disease, one that is solidly grounded in microbiology, a medical science. In this chapter, I will investigate whether a more recent scientific theory can serve to represent diseases of the *new* modern medicine. To do so, rather than taking as my starting point the philosophy of science literature and attempting to fit medical theory into an existing template for a scientific theory, I will start with the germ theory—as a clear and unquestionable example of a theory of disease in medicine—and reconstruct this scientific theory of disease.[1]

Germ theory represents a family of diseases—contagious infections—that characterize a period that I described in Chapter 1 as belonging to an earlier modern medicine. In contrast, noncommunicable diseases are characteristic of the new modern medicine. Along with emerging and re-emerging infections, noncommunicable conditions are commonly regarded as the new epidemics—it is not uncommon to read about epidemics of heart disease, diabetes, or cancer. Noncommunicable conditions include those that are commonly recognized as "chronic diseases." In Chapter 4, I will question whether there could be a theory of chronic diseases given their apparent diversity. In this chapter, I will search for a theory of a family of diseases that is no less important to the story of contemporary medicine, and in many ways is symbolic of the new modern medicine: the cancers. For many, cancers have replaced infections as the disease most dreaded. Over the course of the 1900s, as mortality from various infections declined in the West, the mortality rate

[1] I take this project to be of the kind advocated by Maël Lemoine (2015) under the banner of naturalizing disease.

for cancer rose, driven by an increasing incidence of lung cancer (Siegel et al. 2021). By the end of the century, cancer was the second most common cause of death by disease or family of disease (D. Jones et al. 2012). And unlike chronic noncommunicable diseases, which include conditions as dissimilar in their nature and etiology as asthma and Alzheimer's disease, different cancers might seem to have more in common.

Philosophers, historians, and cancer biologists have identified what some have called the "oncogene paradigm" or "somatic mutation theory" of cancer as an important theoretical landmark in cancer research at the end of the 1900s (Morange 1997; Sonnenschein and Soto 2000; Bertolaso 2016). While throughout the 1900s much cancer research following the blueprint of the germ theory took place, by the end of the century consensus had coalesced around a new theory. The somatic mutation theory (SMT) holds that the hallmarks of the cancer cell are caused by mutation or altered expression of cellular genes, including oncogenes and proto-oncogenes. SMT has an important predecessor in what epidemiologists referred to as the "multistage theory" of cancer (Armitage and Doll 1954; Frank 2007; Plutynski 2013). Multistage theory was an epidemiological theory that accounted for the observation that adult cancers increase in incidence with age by modeling cancer progression as a series of rate-limited stages, commonly proposed by its adherents to be genetic mutations.

Rather than theories, some have thought of germ theory and SMT as scientific frameworks—"paradigms" or "research programs"—because of their role in guiding research by specifying an etiological blueprint for the causes of contagious and cancerous diseases. Either way, SMT has been hugely influential in shaping the way scientists, doctors, and the public think about cancer as well as the landscape of cancer research. With SMT, cancer became a standard-bearer for a new molecular understanding of disease (Morange 1997, 2007). Arguably, the theory has not had the same therapeutic impact as germ theory (at least not yet). However, precision oncology is a recent trend in medical research aiming to harness knowledge of the molecular basis of cancer to improve cancer diagnosis and treatment. (I determine what exactly precision medicine amounts to in Chapter 9.)

Regardless of its therapeutic success, the somatic mutation theory of cancer is a useful case study in concepts and theories of disease at the turn of the twenty-first century, especially given its relationship to molecular biology. When George Engel argued in 1977 that medicine was ruled by a "biomedical model," he declared that molecular biology was its "basic scientific

discipline" (130). Central to Engel's criticism of the biomedical model and to the concerns of other critics following Engel's lead is the model's purported "reductionism": the model in some sense decomposes the patient into biological parts—a bag of molecules, including genes. One problem with this approach in the context of cancer research is the immense complexity of cancer, which might defy simple, bottom-up reductive explanations. Another problem is that by viewing cancer as a disease of cells and genes and molecules, SMT excludes important environmental and social causes. These two problems comprise "molecular reductionism," a criticism that I will evaluate at the end of the chapter.

Modern scientific theories of disease unify models of the pathogenesis of specific kinds of disease within a family. For instance, SMT, a theory of cancer, represents cancer as a family of diseases characterized by uncontrolled cell proliferation, genetic alteration, and molecular dysregulation. In section 2.2, I examine a prototypical modern theory of disease, the germ theory of contagious disease. In section 2.3, I reconstruct SMT in germ theory's image and identify an important epidemiological forerunner to SMT in the multistage theory of carcinogenesis. In section 2.4, I defend my proposal against alternatives that emphasize frameworks of disease rather than theories. Last, in section 2.5, I disambiguate the classical critique of molecular reductionism and evaluate whether it is a fair criticism of SMT.

## 2.2. The Germ Theory, a Prototypical Modern Theory of Disease

During the nineteenth century, waves of cholera and smallpox epidemics battered Western Europe and its colonies, while endemic diseases like pneumonia, diphtheria, tuberculosis, and infectious diarrhea consistently dragged bodies under.[2] For centuries, doctors theorized that such diseases were caused by contagia (seeds or poisons that could be spread from person to person) or by miasmas (deadly vapors rising from filth and decay) or by both. In the middle of the century, a miasmatist zymotic theory held sway. Zymotic theory said that infectious air was produced by a process of chemical decay, which might involve microorganisms but only as a byproduct of

[2] The secondary historical literature on germ theory is vast. In this section, I draw most heavily from Temkin (1977), King (1991), Tomes (1998), Carter (2003), and Snowden (2020).

the chemical reaction. By the end of the century, laboratory research—the science of late nineteenth-century scientific medicine—had turned the tide in favor of germ theory and against the zymotic theory and other rivals.[3]

The zymotic theory had a cousin in chemical theories of alcoholic fermentation and of putrefaction (spoilage) of food. These theories often went hand in hand with the doctrine of spontaneous generation, the theory that the germs found in fermenting or putrefying food could be produced de novo out of the chemical process. In the 1850s and 1860s, Louis Pasteur, a chemist, used exacting experiments to support an alternative germ theory of fermentation and putrefaction: microbial germs were not spontaneously generated in nutritive samples but were deposited from the air and contaminated the samples, causing the subsequent putrefaction. Pasteur drew an analogy between fermentation and human disease, speculating that germs were also responsible for the latter. In 1865, Lister applied the germ theory of putrefaction to the problem of surgical wound infections and proposed his antiseptic solution, making the analogy material. Finally, in the 1860s and 1870s, experimentalists (including Pasteur) used animal experiments to demonstrate that microscopic germs from cases of anthrax, rabies, and diphtheria could be isolated, cultured, and introduced into animals to reproduce the disease. Robert Koch's discovery of the "tubercle bacillus" as a cause of tuberculosis is often held up as the crowning achievement of this research tradition.

From around 1870, when it was articulated (e.g., in Hogg 1870), the germ theory of disease represented a generalization of these results. Though theories of animate contagion are scattered throughout history, the germ theory of the late 1800s was a relatively coherent and widely shared set of beliefs (Tomes 1998). At first a minority view among physicians, by the 1880s it was the topic of several textbooks, and by the 1890s it was being taught in medical schools as orthodoxy. The physician William H. Mays describes the theory well in a paper delivered to the San Francisco Medical Society:

> I hold that every contagious disease is caused by the introduction into the system of a living organism or microzyme, capable of reproducing its kind and minute beyond all reach of sense. I hold that as all life on our planet is

[3] Although John Snow proposed a germ theory of the cause of cholera in his 1855 *On the Mode of Communication of Cholera*, in which he used epidemiological as well as clinical evidence and reasoning, he never visualized the cholera bacterium microscopically and did not manage to convince most of medicine of his theory.

> the result of antecedent life, so is all specific disease the result of antecedent specific disease. I hold that as no germ can originate de novo, neither can a scarlet fever come into existence spontaneously. I hold that as an oak comes from an oak, a grape from a grape, so does a typhoid fever come from a typhoid germ, a diphtheria from a diphtheria germ; and that a scarlatina could no more proceed from a typhoid germ than could a seagull from a pigeon's egg. (1880–1881, 110)

Mays's address contains all the central ideas of germ theory as described by historian Nancy Tomes (1998, 33): "[T]he germ theory consisted of two related propositions: first, that animal and human diseases were caused by distinctive species of microorganisms, which were widely present in the air and water; and second, that these germs could not generate spontaneously, but rather always came from a previous case of exactly the same disease."[4] Tomes notes that not all the theory's proponents saddled it with the exact same commitments. For example, in his book *The Germ Theory Applied to the Explanation of the Phenomena of Disease*, physician Thomas J. MacLagan writes, "The theory is, that many diseases are due to the presence and propagation in the system, of minute organisms having no part or share in its normal economy. This it is, and no more" (1876b, 1).[5] MacLagan was keen to separate the contentious debate over whether life could arise through spontaneous generation (for MacLagan, a biological question) from the matter of the cause of a disease (a question for pathological science). He thought the germ theory ought only answer the latter question (MacLagan 1876a, 1876b).

Nonetheless, among key historical expressions of germ theory, three central principles recur. K. Codell Carter (2003) calls these the Bacterial Hypothesis, the Dissemination Hypothesis, and the Distinguishability Hypothesis. I will refer to Carter's Bacterial Hypothesis as the Germ Hypothesis because by the end of the 1800s it was suspected that germs too small to be visualized with a light microscope (but now visualizable using an electron microscope), such as Pasteur's rabies germ—what are now known as viruses—also cause human disease.[6] We can frame these commitments

[4] Phyllis Richmond (1980) provides a similar characterization of germ theory.

[5] See also C. V. Chapin in *The Present State of the Germ Theory of Disease*: "This theory maintains that many pathological conditions are caused by the presence in the body of certain microorganisms" (1885, 4).

[6] Nineteenth-century germ theorists often used terms like "bacteria" and "viruses" less precisely than doctors do today, to refer to any microscopic germ.

roughly as follows, while recognizing that not every germ theorist associated the germ theory with all three.

**Germ Hypothesis:** certain diseases are caused by living microscopic germs.

**Dissemination Hypothesis:** disease-causing germs do not arise spontaneously in the host but are introduced into the host from without.

**Distinguishability Hypothesis:** a specific kind of disease is caused by the same specific kind of germ and no other ("same cause"); and a specific germ causes the same specific disease and no other ("same effect").

As seen above, germ theory consists of a few general theoretical principles or hypotheses about the pathogenesis of certain diseases or about how certain diseases are caused. More specifically, these hypotheses represent one aspect of pathogenesis: the external etiology or origin of disease (rather than, for instance, internal constitutional factors or the progression of disease).

The Germ Hypothesis is the indispensable core of germ theory to which all these prominent self-avowed germ theorists subscribed, and some of them even associated germ theory with *only* that hypothesis. Whatever one's commitment to the other two hypotheses, there is no *germ* theory of disease without the hypothesis that at least some diseases are caused by germs. However, which diseases? The Germ Hypothesis also sets the scope of the theory, the group of diseases to which it applies. Here there is slight variation among theorists. MacLagan claims that germ theory applies to "many diseases" and goes on to clarify that germ theory is believed to cover contagious and epidemic diseases as well as wound infections. Mays applies germ theory to "every contagious disease."[7] Others state that the theory applies to many diseases and explore its scope through examples that we now recognize as diseases of infectious origin (Gradle 1883; Chapin 1885). But against Carter's (2003) own framing of the Germ Hypothesis, none of these theorists claimed that germ theory applied to *all* diseases.

I also depart from Carter's formulation of the Distinguishability Hypothesis, which he identifies with "same cause" only (as does Tomes in the description of germ theory quoted above). It is sometimes unclear in their written statements whether a germ theorist is endorsing "same cause," "same

[7] Because prion diseases are contagious diseases that are spread by a protein rather than a living microbe, they are contemporary exceptions to the Germ Hypothesis as applied to all contagious diseases.

effect," or both. However, many theorists likely had in mind both senses of specificity. Germ theory was initially suggested—even justified—through botanical analogies (Tomes 1998; Snowden 2020). The germs of disease were much like the germs of a plant: they were living seeds produced in one individual and implanted in another, where they would germinate and produce disease anew. Moreover, "as an oak comes from an oak, a grape from a grape, so does a typhoid fever come from a typhoid germ, a diphtheria from a diphtheria germ" (Mays 1880–1881, 110). This idea is what the Distinguishability Hypothesis calls "same cause." The analogy is naturally extended: just as an acorn always gives rise to an oak tree and never to a maple tree, a typhoid germ always gives rise to a typhoid fever and never to a diphtheria ("same effect").

Mays (1880) compared germ theory to Newton's theory of gravitation. MacLagan called it a "definite theory" in pathology, the science of disease (1876b, 1). Historians have generally concurred with their sources in considering it to be a bona fide theory (Tomes 1998; Carter 2003; Snowden 2020), comparable to stock theories found in physics and biology. For example, Carter (2003) likens the Bacterial Theory of Edwin Klebs (a mid-1870s example of a germ theory) to the ideal gas theory and claims that it is the first explicit theory of disease in the history of medicine (apparently setting aside the example of Hippocratic humoral theory). But should we consider it to be a scientific theory? Does modern medicine (as opposed to ancient medicine, medieval medicine, or another earlier period of Western medicine) really have "theories of disease" in any meaningful sense?

Paul Thompson (2011) answers in the affirmative, arguing that models and theories are common in medical science, from the Bolie (1960) model of insulin-glucose regulation to the Kermack and McKendrick (1927) theory of propagated epidemics.[8] According to Thompson, medical theories and medical models perform similar functions; for example, they both explain and enable control over nature. They do so by virtue of their theoretical content. For instance, germ theory is about the causal etiology or pathogenesis of contagious diseases and can thus be used to explain why people develop diseases that seem to spread from people, animals, or the environment or that occur in clusters or epidemics. But what distinguishes theories from

[8] The Kermack and McKendrick theory is related to the germ theory in that it models the spread of contagion in a way that is consistent with person-to-person patterns of transmission. However, while germ theory represents the spread of contagious disease among individuals, the Kermack and McKendrick theory represents the dynamics of contagious epidemics in populations.

models, according to Thompson, is that theories characteristically unify and integrate a large body of knowledge, including many models, and thus have greater generality. For example, in medicine, an endocrine theory of hormone self-regulation unifies a number of models of more specific hormone systems (Thompson 2011).

If the existence of modern theories of disease is in any doubt, the existence of modern models of disease should not be, if among models of disease we include scientific representations of the pathogenesis or pathophysiology of a particular kind of disease. Open a medical textbook or scientific review article and turn to the section on the pathogenesis of a specific type of disease and you will likely find pictorial and/or linguistic representations of the mechanisms through which the disease is generated (Simon 2008). For example, tuberculosis is caused by a mycobacterium found in the small airways of patients with the disease, which is aerosolized when the sufferer coughs, talks, or breathes and is inhaled into the lungs of its next susceptible victim, where it often remains latent, cordoned within a granuloma by the host's immune cells (Pai et al. 2016).

If we adopt Thompson's reasonable distinction between scientific theories and models, germ theory is a scientific theory because it integrates a large body of knowledge, namely, knowledge about causal etiology spanning most contagious diseases. Germ theory integrates this body of knowledge in the sense that it is supported by a range of data and explanations across several sciences (Plutynski 2013). For instance, in the 1800s it was bolstered by the results of many animal inoculation experiments involving many distinct diseases. But why should we consider germ theory to be a theory *of disease*? After all, unlike theories from the earlier history of medicine, it is not meant to capture the etiology of all diseases. Yet it *does* capture the etiology of a family of diseases—contagious diseases—that share the feature that they are all transmitted from one person to another. It provides further unification to this family of diseases by unifying models of the pathogenesis of particular kinds of contagious disease. In the process, the theory alters the way we think about these diseases: no longer are they merely contagious diseases, but they are also (perhaps primarily) conceptualized as "infectious diseases," diseases produced by infection with a microbe.[9]

[9] That said, it is worth sticking to the idea that germ theory is a theory of contagious diseases rather than considering it a theory of infectious diseases, for two reasons. First, and most important, its original framers, the germ theorists, saw it as a theory of contagious diseases. And second, it turns out that the germ theory does not accurately represent the pathogenesis of all infections. In

A theory of disease, then, unifies a family of diseases by unifying models of their pathogenesis. By "unifying" those models, I mean that the theory performs what we can call "unification by abstraction." A theory of disease unifies models of the pathogenesis of a family of diseases by abstracting away details that are unique to particular kinds of disease, leaving a small number of shared features represented. Germ theory abstracts away details about the kind of germ, its mode of transmission, and its characteristic disease. *M. tuberculosis* becomes a living microscopic germ, its pathogenesis is characterized as a generic process of a germ invading its host's body, and its specific tubercular disease is left unnamed and undescribed. The result is a small number of coarse-grained and highly general hypotheses about the pathogenesis of contagious diseases.[10]

Unification by abstraction differs from the classical account of theoretical unification in philosophy of science: intertheoretic reduction (Nagel 1961). In intertheoretic reduction, less general theories or models are shown to be deducible from a more general theory or model (given "connectability assumptions"). Thus one can derive the less general theories/models from the more general theory. In unification by abstraction, the direction of "derivation" is reversed. The more general theory is "derived" through abstraction from the less general theories/models. In fact, the less general theories/models *cannot* be reproduced from the more general and abstract theory because abstraction involves loss of information. Thus, while it might be thought that in intertheoretic reduction the more general theory is "all the scientist needs" to save the phenomena represented by the less general theories/models, the same cannot be said in unification by abstraction. A theory of disease unifies disease-specific models, but it does not render them as mere deductive implications of a more general and complete theory.

particular, it fails to accurately represent infections caused by microbes that comprise the body's typical microbiome (such as yeast infections).

[10] Paul Thagard (1999) agrees that there is unification across kinds of diseases in medicine. However, he is skeptical that there is a germ theory or a theory of cancer. Thagard instead argues that there is a hierarchy of explanation schemas in medicine. According to Thagard, an explanation schema for a type of disease states the causal factors (such as *M. tuberculosis*) that provide a causal explanation for the disease (tuberculosis) and its symptoms. At a higher level of abstraction, a more general explanation schema unifies all infectious diseases or all cancers. However, in medicine, our theoretical understanding of particular diseases is not exhausted by causal explanations of the kind Thagard has in mind. Rather, it takes the form of more complicated models of disease that allow us to answer a range of questions and explain a range of phenomena, including (for instance) about the mode of transmission of the germ. Moreover, despite Thagard's skepticism, we evidently *do* have a germ theory of disease (or several versions of the germ theory) that was explicitly laid down by nineteenth-century bacteriologists and doctors.

One might worry that germ theory doesn't look very much like theories familiar from physics, such as Newton's laws or quantum field theory, or even from biology, such as Darwinian theory or population genetics; germ theory's generalizations are not universal laws of nature that apply to all objects studied by a particular science. Nor does it resemble either of the structures typically associated with scientific theories by the two historically dominant views on theory structure in philosophy of science, the syntactic and semantic views (Savage 1990); it neither appears to be an axiomatized set of (typically, formal logical) sentences (the syntactic view), nor a set of nonlinguistic (typically, mathematical) models (the sematic view).[11] These two concerns are related because whether philosophers think scientific theories are logical structures or mathematical structures (or something else) is likely influenced by the theories they take to be prototypical. For instance, Thompson (2011) takes medical theories and models to be mathematical, consistent with his support for the semantic view (Thompson 1983). However, the examples that Thompson uses are all dynamical models that can be readily expressed using differential equations. One of the theories of cancer I discuss in the next section, the multistage theory, is similarly mathematical because it is similarly a dynamical model.

Yet many scientific models in medicine are not mathematical. In fact, most models of disease are represented linguistically or pictorially. Mathematical models from biology or physics thus serve as a poor prototype for a theory of disease when it comes to theory structure.[12] Moreover, as Lucie Laplane (2014) argues, the syntactic and semantic views are not meant to identify scientific theories but rather to analyze their structure once theories are identified; therefore, using them to ask whether something is a scientific theory is putting the cart before the theory, so to speak. My main concern in this chapter is not with the structure of theories of disease but with their content and functions. Accordingly, I understand a theory of disease to be any theoretical structure that unifies models of disease pathogenesis, regardless of whether that structure consists of natural language hypotheses, logical formulae, mathematical formulae, or something else.

[11] I will set aside the question of whether we could fruitfully axiomatize the germ theory's hypotheses.

[12] Kenneth Schaffner (1993) argued that a focus on theories like evolutionary theory and population genetics resulted in a sampling bias for philosophers of science studying theory structure in biology. Rather than universal relationships of wide scope, Schaffner argued that the biomedical sciences contain "theories of the middle range."

My approach to characterizing theories of disease is medicine-first. Rather than starting with prototypical scientific theories in science or in philosophy of science and then seeking structures in medicine that resemble these prototypes, I have started with probably the most well-known and influential theory of disease in medicine of the past two centuries, using it as my prototype for a modern theory of disease. I then analyzed its central theoretical features. Philosophers should not decide for doctors whether germ theory is aptly considered a theory of disease any more than they should decide for physicists whether quantum field theory is a physical theory.

Even if we accept that germ theory is a theory of disease, why should we treat it as a prototypical modern theory of disease? My reasons for doing so are historically motivated. First, while it is not the first theory of disease, it is probably the most influential theory of disease of nineteenth-century laboratory science. That is, it is integral to the story of late nineteenth-century modern medicine, which I argued in Chapter 1 was characterized by its close alignment with the theory and methods of the laboratory medical sciences. In the case of germ theory, the relevant science is mainly bacteriology or microbiology (which the theory helped to establish in the nineteenth century). David Wootton (2007) even argues that we should identify the start of modern medicine with germ theory. In investigating theories of cancer, a family of diseases characteristic of the *new* modern medicine, germ theory is thus not only a useful starting point but also an illuminating contrast. Further, germ theory is probably our most well-characterized example of a theory of disease from any period in modern scientific medicine, frequently discussed by historical actors and widely studied by historians.

The germ theory of the late nineteenth century serves as a useful prototype for those interested in modern theories of disease. Guided by its example, we can now scan the field of cancer research to see if there are any theories of cancer in its image, that is, sets of general theoretical structures that unify models of the pathogenesis of specific cancers.

## 2.3. Theories of Cancer

In the nineteenth and twentieth centuries, research on cancer revealed its cellular, genetic, and molecular basis, culminating in a somatic

mutation theory of cancer in the 1980s. SMT views cancer as a disease of dysregulated somatic cell proliferation due to mutations in genes involved in regulatory molecular networks. It remains the dominant theory of carcinogenesis today, complicated but not diminished by research in various quarters. While SMT is often seen primarily as the product of laboratory research in molecular biology, it has an important predecessor in multistage theory, an epidemiological theory supported especially by research in epidemiology.

Cancer is an ancient scourge known to Hippocratic medicine, in which tumors were thought to arise from an accretion of black bile that was in excess. However, its modern understanding begins in the 1800s with microscopic research into histology, the body's tissues, that made use of the new achromatic compound microscope. Pathologists, especially in Germany, developed the hypothesis that cancer cells of a primary tumor as well as those found at distant metastatic sites all arise through proliferation of previous cancer cells and that they could ultimately be traced back through a progressively transforming cell lineage to a normal somatic (bodily) cell. This nineteenth-century theory of cancer's origins is the basis for what I will call the Somatic Cell Hypothesis, one of the core ideas of the twentieth-century SMT: cancer was a disease of abnormal somatic cells proliferating out of control and eventually spreading throughout the body. However, there was no consensus on the cause of the transformation of normal cell to cancerous cell, only a clamor of theories.

During the twentieth century, debates about the causes of cancer continued unabated. Throughout the early and mid-1900s, cancer was variously linked through laboratory and epidemiological research to an array of chemicals, radioactive substances, X-rays, occupational exposures, and consumer products (mostly notably, cigarettes), as well as to genetic mutations. The theory that genetic mutations and other genetic changes cause cancer is often thought to have received its strongest backing in the 1980s through research in molecular biology that carried out laboratory experiments on cells involving cancer-causing genes and then sequenced those genes (Morange 1993, 1997). However, epidemiological research in the 1950s and onward provided earlier support for the theory and produced an important forerunner to the SMT in the multistage theory of carcinogenesis (Plutynski 2013, 2018). In section 2.3.1, I discuss the multistage theory of the mid-1900s, an epidemiological theory. In section 2.3.2, I analyze the later SMT, the now dominant perspective on cancer in medicine.

### 2.3.1 Multistage Theory

In the 1950s, a multistage theory of cancer was proposed that modeled cancer as resulting from a series of rate-limited stages, thus explaining observed patterns of cancer incidence. The history of the multistage theory reveals the often-overlooked role of epidemiology in establishing a genetic basis of cancer. Not only did epidemiology provide data consistent with the hypothesis that genetic mutations cause cancer, but the stages modeled by multistage theory were often hypothesized to be genetic mutations. Multistage theory, an epidemiological theory of disease, is thus an ancestor of SMT.

While the occurrence of cancer is difficult to predict in individuals, there are more predictable patterns of cancer occurrence in populations. In particular, the yearly incidence of cancer, the proportion of people diagnosed with cancer in a year, increases with age for most adult cancers (Figure 2.1).

In 1954, these observations in vital statistics led epidemiologists Peter Armitage and Richard Doll to develop what they called "a multistage theory of carcinogenesis": "[t]he theory that human cancer is the end-result of several successive cellular changes" (9). Building on the modeling of others, especially epidemiological modeling by C. O. Nordling (1953) but also models of cancer in lab animals experimentally exposed to carcinogens, Armitage and Doll showed that a mathematical model in which human cancer occurs after multiple rate-limited steps (in their estimation, six to seven steps) fit the age-related incidence data well for common cancers of the stomach, colon, rectum, and pancreas. In their model, the incidence of cancer was proportional to the age of the population raised to the power of approximately 5 to 6, and the rate of transition from one stage of carcinogenesis to another was assumed to be constant with age.

The age-related incidence data for other common cancers—cancers of the prostate, uterus, breast, ovary, and lung—departed from the predictions of the simpler model in that the rate of transition between stages seemed to vary with age. Yet these anomalies were consistent with the hypothesis that endogenous hormone exposure (namely, to sex hormones) varies throughout life. It was already suspected that hormones were involved in the development of the first four of these cancers. Meanwhile, the data on lung cancer could be explained by the varying exposure of different population cohorts to cigarette smoke: smoking was more common in age cohorts born more recently, speeding the multistage progression of cancer for these younger cohorts. Doll himself had been involved in two of the seminal epidemiological studies

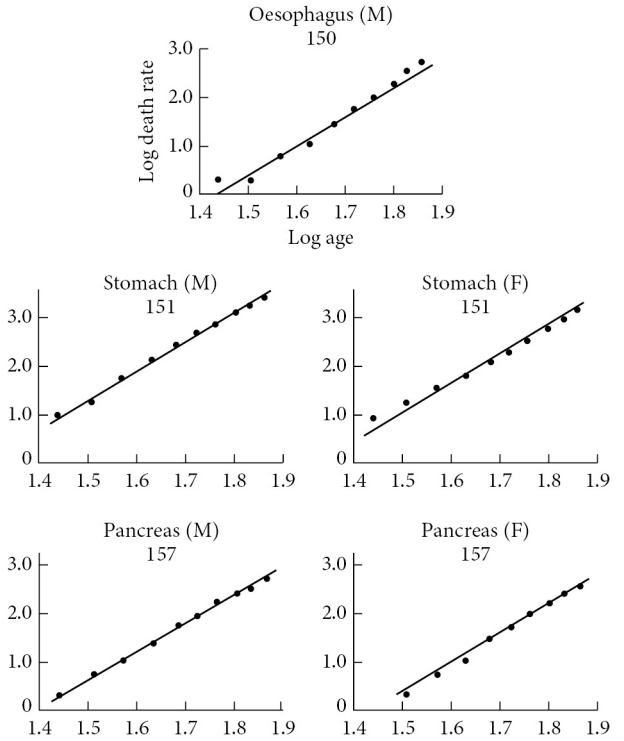



**Figure 2.1** Mortality from cancer versus age for cancers of the esophagus, stomach, and pancreas (in males) and cancers of the stomach and pancreas (in females). Y axis: logarithm of the death rate per million people. X axis: logarithm of the midpoint of the age group.
Armitage and Doll (1954, 3).

implicating smoking as a cause of lung cancer (Doll and Hill 1950, 1954). Thus not only did multistage theory explain why the rate of cancer increased with age, but it also explained the role of suspected carcinogens in the development of cancer: carcinogenic exposures increased the probability that cells would transition through one or more stages in the development of cancer over a given period of time, leading to the occurrence of cancer at earlier ages and a greater cumulative lifetime risk of cancer.

Following C. O. Nordling (1953), Armitage and Doll noted one candidate cause for the cellular transitions in cancer: genetic mutations. Genetic mutation was, after all, thought to be stochastic, explaining the unpredictability of cancer at the individual level. Further, many suspected carcinogens were

thought to be mutagens. Nonetheless, Armitage and Doll (1954) noted that their mathematical analysis did not need to assume any role for genes. In a later review, they concluded that several models and assumptions were consistent with the available data and predicted that evidence to decide among them would likely come either from biology or "quantitative experiments" (Armitage and Doll 1961).

It was quantitative natural experiments in humans (Ashley 1969; Knudson 1971), yielding further epidemiological data, that provided key evidence supporting the role of genetic mutations in the multistage progression of cancer *before* the bulk of the biological research arrived. The most influential study was done by Alfred Knudson (1971). Multistage theory predicted that the age-related incidence curve for populations with fewer stages of cancer to surpass would be left-shifted, resulting in greater cancer incidence earlier in life. If genetic mutations comprised some of the stages, then populations inheriting these mutations rather than acquiring them during life would have fewer stages left to undergo. Knudson tested this hypothesis by comparing the curves for hereditary versus nonhereditary retinoblastoma, a childhood cancer. The incidence curve for hereditary retinoblastoma, retinoblastoma that runs in families, displayed transformations consistent with the removal of one rate-limited step. Hereditary retinoblastoma, like all hereditary diseases, was thought to result from an inherited genetic mutation. Thus Knudson's persuasive results confirmed a role for genetic mutation in the stages of multistage theory.

In particular, the data suggested that retinoblastoma required two mutations or "hits" to occur, and that in familial cases one mutation was inherited while the second was acquired. In 1986, research in molecular biology revealed that both of these mutations occurred in one gene, now known as *RB*, the retinoblastoma gene. *RB* was identified as the first "tumor suppressor gene," one of two kinds of cancer genes in SMT. The initial evidence for the nature of tumor suppressor genes came from Knudson's epidemiological analysis; in their unmutated form they *suppress* carcinogenesis, consistent with the finding that a loss-of-function mutation in both *RB* alleles—two hits—is needed for retinoblastoma to occur.

A version of multistage theory in which the stages are genetic mutations gradually became the leading scientific theory of carcinogenesis (Frank 2007). According to my analysis in section 2.2, multistage theory is a theory of disease because it satisfies several criteria. First, it is a theory in that it explains, and does so by unifying and integrating a large body of scientific

knowledge. This knowledge includes data from sciences such as laboratory biology and epidemiology, as well as more specific models of carcinogenesis, including models that represent particular kinds of cancer. While a two-stage model best explains the development of retinoblastoma, models including six or more stages are generally needed for adult cancers. Multistage theory unifies these models by abstracting away details such as the tissue-specific number of stages and rate of mutation.

Multistage theory is, in particular, a theory *of disease* because the models that it unifies are models of the pathogenesis of particular kinds of disease comprising a family of diseases (cancer). In particular, multistage theory represents carcinogenesis as a multistage process and, when linked to the hypothesis of genetic mutation, explains cancer occurrence as caused by a series of genetic changes to somatic cells. Anya Plutynski (2013, 2018) similarly argues that multistage theory integrates data and explanations about cancer from multiple areas of science and unifies various cancers as of a kind, calling it "a unifying explanatory theory of carcinogenesis" (2013, 472).

Plutynski also rightly notes that the primary explanatory target of multistage theory is patterns of cancer incidence instead of individual cases of cancer. That makes multistage theory an "epidemiological theory," according to epidemiologist Nancy Krieger's (2011) definition. Given that epidemiology is typically defined as the study of the distribution and determinants of health and disease in populations, Krieger defines epidemiological theories as "*theories of disease distribution*" in populations (29). Krieger distinguishes epidemiological theories from what she calls "theories of disease causation," which she suggests theorize about disease mechanisms in individuals. Presumably, the germ theory discussed in section 2.2 and the SMT that will be discussed in section 2.3.2 are good examples of theories of disease causation.[13] Krieger emphasizes that epidemiological theories are "*not* the same as theories of disease causation" (30).

Krieger's contrast calls to mind another distinction, this one proposed by the epidemiologist Geoffrey Rose (2001, 2008), between "the causes of cases" and "the causes of incidence." Rose claims that the causes of differences in cases within a population are typically distinct from the causes of differences

[13] Krieger (2011) suggests that causal pie models in epidemiology are also theories of disease causation. I will discuss causal pie models and their associated "sufficient-component cause theory" in Chapter 3. Meanwhile, a good example of an epidemiological theory is William Kermack and Anderson McKendrick's (1927) "mathematical theory of epidemics," an early example of compartment modeling. Today compartment modeling is used both in infectious disease modeling as well as in multistage theory (Frank 2007).

in incidence between populations. (Interestingly, as an example of the former Rose mentions genetic causes.) In his Intersalt study of thirty-two countries (Intersalt Cooperative Research Group 1988), Rose found that the curve representing the distribution of blood pressure values in a population had a similar shape across countries but varied in its placement along the horizontal blood pressure axis, and thus the mean blood pressure varied too. Rose suggested that these horizontal shifts in population distribution should be attributed to characteristics of populations rather than characteristics of individuals. Similarly, mathematical biologist Steven Frank (2007) argues that multistage theory accounts for shifts in the age-related cancer incidence curve seen, for instance, in Knudson's comparative data on retinoblastoma.

Rose makes it seem as if the causes of changes in the incidence curve are not typically the causes of the cases, while Krieger suggests that theories that account for changes in the incidence curve do not account for cases. Thus it sounds as though we should sharply distinguish multistage theory and the causes it represents from SMT and the causes it represents; in which case, my claim that multistage theory is a predecessor of or forerunner to SMT might be too strong.

However, I've argued elsewhere that Rose's distinction is best seen not as distinguishing different causes operating at the population and individual levels, but instead as distinguishing different (contrastive causal) explanations at the population and individual levels (Fuller 2022a). Multistage theory explains differences in the incidence of, for instance, hereditary versus nonhereditary retinoblastoma by citing differences in the number of mutations that cells need to accumulate—that is, mutations to cells in individuals, causing cancer in those individuals. It explains differences in the incidence of lung cancer in younger versus older age cohorts by citing differences in smoking habits—again, in individuals—that accelerate mutation rates in populations. What distinguishes multistage theory from the later SMT is that while SMT explains why one individual develops cancer while another does not by citing differences in their individual "exposure" to an inherited or acquired mutation, multistage theory explains why the incidence of cancer in one population is greater than in another's by citing differences in population mutation *rates*. As an epidemiological theory, multistage theory offers epidemiological explanations involving the incidences of disease and a population's exposure to causes, but they are *the same causes* that figure in the explanations offered by its more recent cousin, the SMT. Multistage theory offers a "population perspective"

(Fuller 2022a) on the genetic causes of cancer that complements SMT's individual perspective.

More specifically, multistage theory supported a primordial version of the hypothesis that genetic changes cause cancer. This idea was further refined in the 1980s with the discovery of the kinds of genes involved (proto-oncogenes and tumor suppressor genes), at which point the hypothesis evolved to become what I will call the Oncogene Hypothesis. Moreover, multistage theory hypothesized that multiple mutations accumulated within a cell or cell lineage in a process of multistage progression toward cancer. The hypothesis of multistage progression was further confirmed in the 1980s and 1990s by findings in molecular biology and tumor histology and is a core implication of the SMT (R. Weinberg 2014). In summary, epidemiology supplied not only evidence favoring the role of genetic mutations in cancer but also a theory of disease in which those mutations explain aspects of cancer incidence within and across populations. While adopting a different (population-level) perspective on cancer, it is nonetheless aptly considered a forerunner to SMT.

Before turning to SMT, two further implications of multistage theory are worth mentioning. The age-related incidence curves for heart disease and other leading noncommunicable causes of death display similar features as the curves for adult cancers, increasing in slope with age. Multistage theory inspires a general (albeit, more speculative) explanation for these phenomena, the explanation that they result from failure in multiple systems or components. This raises the possibility that a "multicomponent theory" of disease unifies a wide range of noncommunicable diseases, with the biological mechanisms awaiting further elucidation (Frank 2007).

Finally, the multistage theory (and possibly also a multicomponent theory) goes some way toward explaining the epidemic of cancer in the new modern medicine, why deaths from cancer (along with deaths from heart disease and other noncommunicable diseases) have displaced deaths from infectious diseases as the leading causes of death in the developed world. In Chapter 1, this development was introduced as the classic "epidemiologic transition." Part of the explanation for the transition is given by Abdel Omran's (2005) "theory of the epidemiology of population change." As mortality from infections, particularly infections in early life, declined over the decades prior to and including the early 1900s, populations aged. People then began dying of age-related diseases—the diseases just mentioned that increase in incidence as a function of age—at a greater rate. As a theory that

explains changes in the incidence of death and disease, Omran's is another example of an epidemiological theory.

However, while Omran's theory explains why populations aged, it does not explain why, as they aged, they developed cancer (and other noncommunicable diseases) and died from cancer at an increasing rate. Cancer deaths increased not only as a proportion of total deaths, as would be expected as deaths from other causes decreased, but also in absolute incidence, particularly with age. Why did populations acquire more cancer (and other noncommunicable diseases) as they aged? Multistage theory (and possibly a more generalized multicomponent theory) provides an explanation: cancer results from multiple mutations, and consequently cancer and deaths from cancer occur more frequently in older age groups that have had more time to accumulate mutations (while other noncommunicable diseases plausibly result from multiple component failures that accumulate over time).

To the hypothesis that cancer originates in a somatic cell, multistage theory added the hypothesis that carcinogenesis results from genetic mutations. The multistage theory provided an epidemiological perspective on the genetic causes of cancer, explaining epidemiological data on the age-related incidence of cancer and the rising overall incidence of cancer in human populations.

### 2.3.2 Somatic Mutation Theory

The hypothesis that genetic changes cause cancer was further confirmed in the 1980s by laboratory research in molecular biology, and the nature of the genes and their cellular roles was discovered. The result was a coherent theory of carcinogenesis in individuals, the SMT, which explains the dysregulated proliferation of somatic cells in cancer as due to mutations in genes involved in regulatory molecular networks.[14]

A competing strain of twentieth-century research investigated whether cancer was caused by a microbe, whether germ theory and the laboratory techniques that developed alongside the theory could be extended to cancers.

[14] In reconstructing the history of SMT in the 1800s and 1900s, I rely especially on historical work by Victor Triolo (1964), Michel Morange (1993, 1997, 2011), Ton van Helvoort (1999), and Gregory Morgan (2022), as well as the secondary scientific literature (R. Weinberg 1988; Bishop 1995).

In the late 1800s, scientists failed to culture any cancer-causing microbes. In the early 1900s, they finally succeeded in purifying invisible agents that were demonstrated to be the cause of leukemias and other cancers in laboratory animals. In 1911, Peyton Rous published work detailing his experiments on spindle cell sarcoma in chickens. In an inoculation experiment not dissimilar to the ones performed by Pasteur, Koch, and others decades earlier, Rous ground up the cancerous muscle tumors from afflicted chickens, passed them through a filter, collected the cell-free filtrate, and inoculated new chickens with the filtrate, producing histologically similar tumors. The causative agent was later named the Rous sarcoma virus. Rous went on to be a standard-bearer for the hypothesis that viruses were the cause of cancer.

Research on human cancer-causing viruses became more mainstream in the 1960s. Under the National Cancer Plan of 1971, the U.S. established the Special Virus Cancer Program. While the program did not live up to the hopes of identifying viral causes for most human cancers, tumor virus research did yield breakthroughs in molecular biology, including the molecular biology of cancer. In 1976, Michael Bishop and Harold Varmus identified the gene responsible for cancer in the Rous sarcoma virus: *v-src*. They discovered that it was similar in sequence to a gene present in the DNA of non-infected, non-transformed avian cells that they labeled *c-src*. They hypothesized that *v-src* was a mutant form of *c-src* that had found its way into Rous sarcoma virus when *c-src* was accidentally incorporated into the viral genome of a distant viral ancestor and subsequently mutated into *v-src*. They thus called the cancer-causing gene, *v-src*, an "oncogene" and the unmutated gene, *c-src*, a "proto-oncogene."

Michel Morange (1993) argues that scientific consensus about the privileged role of genes in cancer was soon solidified due to research between 1981 and 1984 coming out of the laboratories of the molecular biologist Robert Weinberg, among others. Using a novel "transfection" protocol, Weinberg and collaborators showed that digested DNA from virally transformed cells, chemically transformed cells, and biopsied cells from spontaneously arising human cancer could all transform normal cells into cancer cells in culture. It thus seemed that the same oncogenes found in tumor viruses could also occur spontaneously in humans or through carcinogen-induced mutation and, moreover, from a common set of proto-oncogenes. In some cases, an oncogene was found to differ from its proto-oncogene by only a single DNA base pair substitution, resulting in a single amino acid substitution in the protein encoded. Other oncogenes were found to be switched on through

various other mutational mechanisms—from chromosomal translocations to multiple gene repeats—resulting either in expression of an abnormal cellular protein or overexpression of a normal one. In 1986, Weinberg's group isolated the *RB* tumor suppressor gene that was mutated in retinoblastoma.

While early experiments successfully induced the formation of cancer cells by transfecting a previously immortalized cell line with a single oncogene, Weinberg soon found that non-immortalized cells could be transformed only when they were transfected simultaneously with two oncogenes, suggesting that multiple mutations are required to successively transform a cancer cell lineage from a normal somatic cell to a tumor-forming cancer cell. This laboratory finding provided biological confirmation for the multistage theory of carcinogenesis. But how did these genes produce the malignant cell phenotype?

As the transforming oncogenes were being sequenced by molecular biologists, so too were major molecular players in normal cell signaling. In the mid-1980s, several oncogenes were found to be similar to genes for growth factors, growth factor receptors, and other signaling molecules within the cell membrane and inside the cell. These proteins were involved in cell signaling pathways regulating cell division. That explained why oncogenes encoding abnormal or abnormally expressed versions of these proteins led to dysregulated, hyperactive cell division. For example, it was discovered that *RB* encodes a protein that normally prevents progression through the cell cycle and thus halts cell division. Mutations in *RB* make cells insensitive to antigrowth signals and allow them to cycle and divide uncontrollably. Cancer was a disease of cell proliferation and genetic mutation, but also one of disrupted cell signaling. The cellular vision of cancer from the 1800s had been assigned a genetic and molecular etiology by the end of the 1900s.

In 2000, Weinberg published a landmark review article in the journal *Cell* with Douglas Hanahan titled "The Hallmarks of Cancer." The article consolidates the principles of cancer biology developed over the preceding decades and contains the core propositions of what has been called the somatic mutation theory of cancer, the dominant perspective on cancer at the end of the twentieth century as well as today. The SMT remains under development, so to characterize the theory without too much ambiguity it is useful to settle on a particular time slice in its evolution. Because my reference point for understanding the new modern medicine is the turn of the twenty-first century, I will focus on the SMT as it is reflected in the year 2000 in Hanahan and Weinberg's well-cited paper. It is also useful to focus on a

particular formulation of the theory (a canonical one) because there might be room for disagreement on what the core principles of the theory are. The theory has in some ways expanded since 2000 to accommodate growing research on the immune system in cancer as well as the tumor microenvironment (Hanahan and Weinberg 2011). However, its central hypotheses remained essentially unchanged.

Hanahan and Weinberg (2000) argue that several traits or capabilities—what they call "the hallmarks of cancer"—must be acquired by a cell en route to becoming a cancer cell. Their paper largely enumerates these capabilities and explores their molecular basis in the cell. They are "self-sufficiency in growth signals, insensitivity to growth-inhibitory (antigrowth) signals, evasion of programmed cell death (apoptosis), limitless replicative potential, sustained angiogenesis, and tissue invasion and metastasis" (57).[15] Hanahan and Weinberg propose that "these six capabilities are shared in common by *most and perhaps all types of human tumors*" or "*human cancer*" (57, my emphasis). Together they allow a cell to exhibit unrestrained survival and replication, to recruit a blood supply, and to invade and metastasize to other tissues.

The authors review research on some of the molecular mechanisms responsible for the hallmarks, which continue to be a subject of much research in cancer biology. These pathological mechanisms are conceptualized as malfunctions in the "integrated circuit of the cell" (Figure 2.2).

The integrated circuit consists of a network of signaling pathways regulating cell proliferation and motility. The hallmarks of cancer thus represent dysregulation of cell proliferation and motility due to disruptions in regulatory pathways in the cell (Hanahan and Weinberg 2000, 57). What initiates these defects are typically mutations in proto-oncogenes and tumor suppressor genes (57).[16] Although the biological reality is complex, Hanahan

[15] They list genome instability separately as an "enabling characteristic" because it enables the development of the other capabilities. In their 2011 follow-up paper "Hallmarks of Cancer: The Next Generation," they add "two emerging hallmarks of potential generality to this list—reprogramming of energy metabolism and evading immune destruction" (646) as well as one further enabling characteristic: tumor-promoting inflammation. In a further follow-up published in 2022, "Hallmarks of Cancer: New Dimensions," Hanahan promoted these two "emerging hallmarks" to the status of "core hallmarks" and proposed one further prospective hallmark, phenotypic plasticity, as well as two prospective enabling characteristics, epigenetic reprogramming and polymorphic microbiomes.

[16] Hanahan and Weinberg (2011) gesture toward additional roles for epigenetic changes, or alterations to chromosomes such as methylation that do not change the sequence of DNA base pairs, while Hanahan (2022) highlights them as a potential enabling characteristic. Thus something like "somatic genome alteration theory" would be a better name for their 2011 theory, but I stick with "somatic *mutation* theory" both for the sake of consistency with the philosophical literature and because genetic mutations were emphasized in their 2000 paper, which is my reference point.

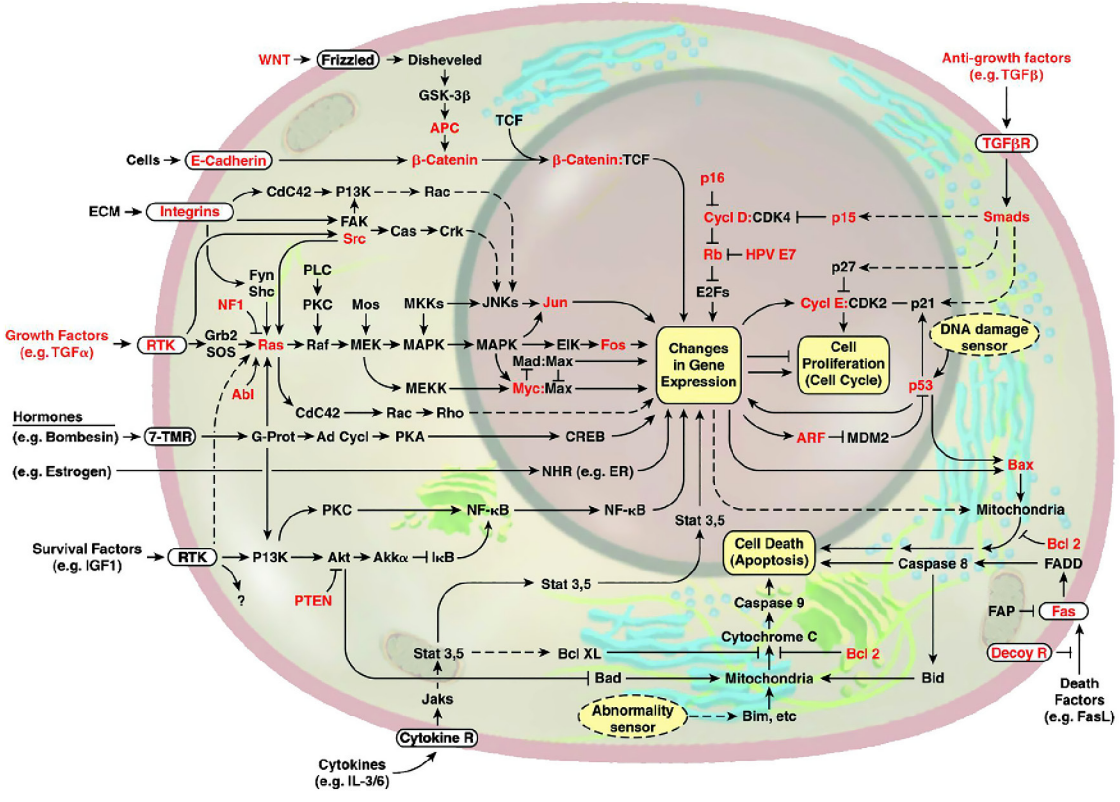



**Figure 2.2** The integrated circuit of the cell.
Hanahan and Weinberg (2000, 59).

and Weinberg present a set of general principles that different cancers obey, while recognizing that mechanisms and mutations responsible for the cancer hallmarks can vary from one type of cancer to another, and even among cases of the same type of cancer. Cancer hallmarks are usually (but not always) acquired in different mutational stages because different genes are usually (but not always) implicated in the acquisition of different cancer hallmarks (66).

The historian Michel Morange (1997) helpfully summarizes the main components of SMT that emerged in the 1980s and that were later articulated in their more mature form by Hanahan and Weinberg in their 2000 review, referring to them as "the oncogene paradigm."[17]

I will instead refer to the principles articulated by Hanahan and Weinberg (2000) as the somatic mutation theory because, as I will argue, they constitute a theory of cancer as a family of diseases. A "somatic mutation theory" has often been discussed by philosophers (Marcum 2005; Malaterre 2007; Bertolaso 2016) following the use of this expression by cancer biologists Carlos Sonnenschein and Ana Soto (1998, 2000) in reference to the consensus captured in articles such as "The Hallmarks of Cancer." Some of the discussions about SMT have been critical. My purpose in this section and in the next is not to evaluate the adequacy of SMT as a theory but to reconstruct it and understand its role in cancer research and medicine. In section 2.5, I evaluate an important criticism of SMT, its purported "reductionism."

As we did with germ theory, we can parse SMT roughly into three hypotheses found in Hanahan and Weinberg's (2000) paper:

**Somatic Cell Hypothesis:** cancer arises through proliferation of cancer cells that descend from a previously non-cancerous somatic cell of a specific tissue and cell type.

**Oncogene Hypothesis:** the hallmarks of a cancer cell are caused by mutation or altered expression of cellular genes ("proto-oncogenes," "tumor suppressor genes").

[17] "In this paradigm, cancer is due to the increased expression *or* structural modification of a *limited number* of genes called the oncogenes. These oncogenes can be activated by different mechanisms. . . . Cancer is due to a multiple alteration of oncogenes, and therefore to the cooperative action of different oncogenes. Oncogenes code for proteins involved in the control of cell division. . . . (I have not mentioned in this description the antioncogenes or tumor suppressor genes; their discovery came later, and only slightly altered the main components of this oncogene paradigm.)" (Morange 1997, 7). Morange argues that the oncogene paradigm replaced a different regulatory theory of cancer that was dominant between 1960 and 1980 in which cancer is caused by mutations in regulatory sequences within the genome, producing dysregulated expression of other sequences coding for functional proteins.

**Cancer Hallmarks Hypothesis:** these cellular genes encode proteins involved in molecular networks regulating cell proliferation and motility; through disruptions in those networks, cancer cells acquire multiple capabilities ("hallmarks") enabling uncontrolled proliferation and invasion/metastasis.

These hypotheses represent part of the pathogenesis of cancer, specifically how cancer cells arise (carcinogenesis) and the cell-intrinsic causes of this process.

I divide the theory into separate hypotheses partly for convenience but also for several other reasons. Different hypotheses became widely accepted (or at least widely entertained) at different times: the Somatic Cell Hypothesis in the 1800s (Rather 1978), the Oncogene Hypothesis between the 1950s (Armitage and Doll 1961) and the 1980s (Morange 1997), and the Cancer Hallmarks Hypothesis by around 2000 (Hanahan and Weinberg 2000). The hypotheses are to some extent independent, and one could accept or entertain one hypothesis but reject or set aside another. For instance, through most of the 1900s, one might accept the Somatic Cell Hypothesis along with a "viral hypothesis" stating that the cancer cell phenotype was due to infection with a tumor virus, which differs from the Oncogene Hypothesis that eventually surpassed the viral one. Moreover, different scientists, historians, and philosophers have placed different emphases on these hypotheses when articulating SMT.[18] Whereas the Germ Hypothesis was the central component of germ theory to which prominent germ theorists subscribed and without which there is no sense in which the theory is a *germ* theory, the Oncogene Hypothesis is the minimal commitment of the SMT that is most regularly associated with the theory and that justifies calling it a theory of somatic *mutation* or an *oncogene* paradigm. However, I have included all three hypotheses in my reconstruction of SMT because all three ideas appear in Hanahan and Weinberg (2000) and recur throughout the historical and philosophical literature on SMT.

The Somatic Cell Hypothesis sets the scope of the theory: it is a theory about how "cancer arises," and it applies to "most and perhaps all types of

[18] According to Marta Bertolaso, "The clonality of tumor cells and the determinant role of somatic mutations are the fundamental tenets of the SMT" (2016, 19), which partially captures the Somatic Cell and Oncogene hypotheses. Plutynski associates SMT with what I'm calling the Oncogene Hypothesis: "The somatic mutation theory is the view that cancer is essentially the product of the gradual accumulation of cell-intrinsic changes: mutations, during somatic cell division" (2018, 201).

human tumors" or "human cancer" (Hanahan and Weinberg 2000, 57), that is, solid tumors as well as blood neoplasms. (Oncologists typically add that cancerous tumors are able to spread to other tissues in order to distinguish them from non-malignant tumors.) While this hypothesis locates the origin of a cancer cell lineage in a previously non-cancerous somatic cell, some cancers arise from germ cells, which thus provide exceptions, albeit uncommon ones. The assumption that cancers arise from a *single* somatic cell (and that cancer cells are monoclonal, linked through common descent from that cell) is commonly associated with SMT (R. Weinberg 1999; Bertolaso 2016). However, it should not be confused with the idea that *tumors* are monoclonal, consisting only in cancer cells linked through common descent. Hanahan and Weinberg (2000) recognize that tumors are heterotypic, consisting not only in cancer cells but also other cells of various types that have been "conscripted" by cancer cells through intercellular signaling.

The Oncogene Hypothesis is about the causes that initiate cancers. While Rous thought these might turn out to be yet undiscovered human tumor viruses (van Helvoort 1999), molecular biologists in the 1970s and 1980s determined that the common factor among cancers was changes to the genome that could be brought about by a range of external mutagens, including viruses but also chemicals or radiation, or they could be inherited or occur spontaneously. The genetic mechanisms producing oncogene expression are also varied. Cancer-causing genes are thus a bottleneck through which diverse etiologic factors exert their effect on cell proliferation. This picture now appears simplistic, given that tumor promoters from alcohol to chronic infection with hepatitis viruses may cause cancer not through direct mutation but rather by promoting inflammation and cell proliferation/turnover (Coussens and Werb 2002; Fujiki et al. 2013). However, strictly speaking, these complications may not contradict the Oncogene Hypothesis, given that mutations or epigenetic changes affecting gene expression may still lie downstream from these mechanisms (Hanahan and Weinberg 2011). While Hanahan and Weinberg (2000) note that genetic changes and other cell-intrinsic causes are not sufficient for cancer progression and that the tumor microenvironment also plays a causative role, somatic mutation theorists like Robert Weinberg (1988) and Michael Bishop (1995) have sometimes suggested that genetic changes are *always* among the causes of a cancer.[19]

[19] In an article titled "The Genetic Origins of Human Cancer," Robert Weinberg (1988, 1966) makes something resembling a necessity claim: "Only after the requisite complement of oncogenes and analogously acting genes have been activated will the cell reach the phenotype

SMT contains nothing resembling germ theory's Distinguishability Hypothesis. For instance, in SMT specific types of cancer are not caused by specific carcinogens or by specific genes, which do not exhibit a one-to-one relationship with recognized types of cancer. In practice, different kinds of cancer are normally distinguished according to cell and tissue type (e.g., small cell lung cancer, basal cell skin cancer) and only sometimes are subtypes distinguished by genotype or gene expression (e.g., breast cancer in which the HER2 protein is overexpressed). In principle, we could modify the Somatic Cell Hypothesis to include the idea that specific kinds of cancer arise from specific tissue and cell types, and that specific tissue and cell types give rise to specific kinds of cancer. However, this principle is typically not associated with SMT, perhaps because it is a taxonomic rather than an empirical principle (more on this point in Chapter 3).

Finally, the Cancer Hallmarks Hypothesis broadens the molecular focus of SMT from genes to regulatory networks controlling proliferation of the cancer cell (Bechtel 2018). The molecular unit associated with a cancer hallmark is not necessarily an altered gene but rather an altered molecular pathway. Again, the theory simplifies the reality of cancer by leaving out the tumor microenvironment beyond the cancer cell as well as the immune system.[20] But once more this omission may be better viewed as an abstraction away from (relevant) detail rather than an idealization that introduces falsity into the theory. In fact, features outside the cancer cell may interact with molecular pathways within the cell, as when cancer cells secrete signaling molecules that recruit stromal cells (Lüönd et al. 2021).

Consistent with multistage theory, one could also include in SMT the hypothesis that cancer progression is a multistage process, in which new cancer cell sub-lineages emerge over time, each sub-lineage acquiring new genetic

of a truly malignant cell." In an article titled "Cancer: The Rise of the Genetic Paradigm," Bishop generalizes, "It now appears that most if not all forms of cancer may contain damage to one or another specific proto-oncogene" (1995, 1313). Plutynski (2020) expresses skepticism that there are any theories of cancer that make lawlike generalizations about necessary and sufficient causes of cancer. However, the Oncogene Hypothesis as I understand it is not a claim about causal necessity or sufficiency but rather the claim that cancer in humans is generally (always or usually) caused by genetic or epigenetic changes, which is not the denial that there are other relevant causes nor the assertion that cancer could not possibly occur without genetic/epigenetic causes. On the separate question of whether this central role for genes in the development of cancer makes cancer a "genetic disease," see Plutynski (2018).

[20] The hallmarks of cancer are discussed by Hanahan and Weinberg as if they are capabilities intrinsic to cancer cells. Bertolaso (2016), Laplane (2016), and Plutynski (2018) all advance or consider more relational views of cancer, cancer hallmarks, or cancer stem cells.

alterations enabling the emergence of new cancer hallmarks that "collaborate" with one other to produce the malignant phenotype. The multistage progression of most cancers may follow from the fact that somatic cells generally lack the hallmark capabilities of a malignant cancer cell, as well as from the (partial) modularity of molecular pathways regulating distinct cancer hallmarks, so that most somatic cells must acquire mutations in multiple independent pathways to become fully malignant. Thus, rather than an independent hypothesis, the multistage progression of cancer is a consequence of the three hypotheses I identified, plus some assumptions about the typical capabilities of somatic cells and the modularity of regulatory pathways implicated in carcinogenesis.

What makes SMT a theory is the same characteristic that made germ theory and multistage theory "theories": SMT has theoretical or explanatory content, and it unifies and integrates a body of knowledge. SMT is about the causal etiology of cancer, how cancer cells arise, and thus can explain why and how malignant tumors develop; it has theoretical and explanatory content. Moreover, SMT is a theory rather than a model because it integrates a diverse body of findings, namely, about the epidemiology, histology, and molecular biology of cancer. For instance, it reconciled research on diverse carcinogens from chemicals and X-rays to viruses, including the results of laboratory experiments and epidemiologic observations on the age-related incidence of cancer. According to the central characteristic of scientific theories discussed in section 2.2, because SMT integrates a diverse body of data and explanations, it deserves theory-level recognition.

As with germ theory and multistage theory, what makes SMT a theory *of disease* is that it unifies models of the pathogenesis of particular kinds of disease that together constitute a family of diseases. As with multistage theory, the family that SMT represents contains diseases characterized by tumor growth and invasion/metastasis, namely, the cancers. SMT unifies models of the pathogenesis of these diseases through "unification by abstraction," by omitting details about the specific tissue and cell types, genes, and molecular pathways implicated in different types of cancer. SMT thus unifies diseases as diverse as pediatric retinoblastoma and adult colorectal cancer (adenocarcinoma), diseases that strike at different ages, affect different parts of the body, and have different clinical courses (Dimaras et al. 2015; Kuipers et al. 2015).

SMT captures what turn out to be considerable similarities between these two otherwise dissimilar diseases. Each of these cancers is typically associated with an initial driver mutation in a tumor suppressor gene (*RB* in

retinoblastoma, *APC* in colorectal cancer). When loss-of-function mutations occur in both copies of *RB* in the genome of a retinal precursor cell, it initiates carcinogenesis in the retina; when mutations occur in both copies of *APC* in the genome of an epithelial precursor cell in the colon, it initiates carcinogenesis there. Why? Because these genes are involved in molecular networks regulating cell proliferation: *RB* blocks cell cycle progression, while *APC* promotes loss of stem cell–like proliferative capacity (among other effects). Unrestrained cell cycling and stimulation of growth receptor pathways are among the hallmarks of cancer, and they result in the uncontrolled proliferation that defines cancers. SMT views retinoblastoma and colorectal cancer abstractly as two examples of a general process of carcinogenesis characterized by genetic alteration, molecular dysregulation, and uncontrolled proliferation. SMT transformed medicine's understanding of what cancer is: cancers are no longer exclusively viewed simply as macroscopic tumors expanding and invading but also as microscopic cells transforming genetically and molecularly.

In summary, just like germ theory, SMT unifies a family of diseases. It also takes after germ theory in employing only a small number of basic principles that bring order and general understanding to a field of medicine. To quote Hanahan and Weinberg again, SMT is "an organizing principle that provides a logical framework for understanding the remarkable diversity of neoplastic diseases" (2011, 646).[21] Germ theory and SMT are rarely explicitly discussed or mentioned in medicine today and are even more rarely recognized by name. As I will discuss in the next section, they are not commonly called upon to do direct explanatory work, which goes some way to accounting for their obscurity in medicine. Moreover, the hypotheses they comprise and the ideas about disease that they encode are so commonplace and accepted that they rarely merit discussion, except when they are being criticized or defended, or in a textbook or a bird's-eye view of a field of research. (These are some of the sources from which I have drawn in reconstructing both

[21] Harold Kincaid (2008) argues that Weinberg's idea that cancer has a logical and coherent theory is not true because it simplifies the complexity of cancer research; rather, suggests Kinkaid, cancer research seeks and offers piecemeal causal explanations, especially those that suggest ways to intervene. On Kincaid's latter point, to accept that cancer research has a theory of cancer is not to deny that it also has and seeks piecemeal causal knowledge. But partly what justifies considering piecemeal explanations about cancer to be part of a body of research (cancer research) rather than an unrelated collection of causal facts that do not comprise a body of research is that there are generalizations that unify and integrate these piecemeal causal facts. On the former point that a theory of cancer simplifies the complexity of cancer research, simplification by idealization or abstraction is common in theorizing and modeling, so this is no argument against cancer researchers having a theory.

theories.) Otherwise, their hypotheses lurk in the shadows as background theory, assumptions upon which medical research and practice often implicitly rely.

The Oncogene Hypothesis was widely adopted only after it was shown to explain the remarkable diversity of cancers better than the Germ Hypothesis. While a viral etiology could explain the origin of some human cancers, the Oncogene Hypothesis could explain the origin of most if not all cancers. The resulting division of explanatory labor allocates contagious diseases to germ theory and diseases of uncontrolled somatic cell proliferation to SMT. Consequently, diseases of uncontrolled somatic cell proliferation that are also contagious fall within the scope of both theories and, indeed, are well explained by both. Notably, most cases of liver cancer and cervical cancer are caused by viruses (hepatitis B/C and human papillomavirus, respectively) that generate tumors. After consensus was built around the Oncogene Hypothesis, germ theory was not thereby rejected as a theory of disease but merely restricted in its scope (most cancers were excluded), illustrating that the dynamics of theory change with respect to theories of disease resists a simple model of winner-take-all competition. That is because modern theories of disease (unlike certain premodern theories, such as humoral theory) do not aspire to explain all diseases but only a family of diseases that share important features.[22] In the very least, I have argued that theories of contagion and cancer have this aspiration.

Lucie Laplane (2016, 2018) reconstructs and defends another "theory of cancer" that she calls Cancer Stem Cell (CSC) theory. Since the beginning of the 2000s, CSCs—cancer cells that resemble the body's normal stem cells in their ability to self-renew and differentiate to generate all the phenotypically diverse cells of a tissue—have been heavily researched (Reya et al. 2001; Kreso and Dick 2014). As such, they represent a development in cancer biology largely postdating my index year of 2000. However, it is worth briefly wondering about the relationship between SMT and CSC theory. According to Laplane (2016), CSC theory consists of two hypotheses and three models. The first hypothesis is that CSCs exist (with the capability to self-renew and differentiate, which is known as "stemness"). The second hypothesis is that not all cancer cells are CSCs; in fact, the majority of them are non-CSCs in that they lack stemness. CSC theory's models are a "CSC model

[22] Similarly, Gillies (2005) argues that medicine has multiple theories of disease that are specific to classes of disease and that this explains unique features of theory change in medicine.

of carcinogenesis," a model of resistance to traditional cancer therapy and cancer relapse, and a model of CSC therapy. I will focus on the CSC model of carcinogenesis because it is most comparable to SMT in that it represents cancer pathogenesis.

According to Laplane, in the CSC model, only CSCs can form new tumors (by self-renewing and differentiating), either at metastatic sites or after tumor debulking at the original site; other cancer cells (the majority) cannot. Laplane contrasts the CSC model with a "classical model." In the classical model, most cancer cells can form tumors because they are capable of self-renewal. (There is no special CSC population uniquely endowed with this ability.) In this contrast, the CSC model and the classical model are clearly in competition. However, we should not be too quick to identify the classical model with SMT. At first glance, they endorse different kinds of claims: the former is about the origins and maintenance of tumors; the latter is about the origins and maintenance of cancer cells. On closer inspection, that is because they represent different aspects of the pathogenesis of cancer. We can draw upon an old distinction (not always made) between *carcinogenesis* as the pathogenesis of cancer and *tumorigenesis* as the formation of a tumor (Rather 1978). SMT represents carcinogenesis, especially the process by which tumor-forming cancer cells arise. On the other hand, the classical and CSC models described by Laplane represent tumorigenesis, the process by which tumors are formed from cancer cells.

Therefore, in examining the relationship between SMT and CSC theory, several possibilities remain live, including competition, additivity, or complementarity (cf. Fagan 2017). The two theories are in competition if they are in direct conflict, if they are mutually incompatible. (Even then, multiple results are possible, from "winner-take-all," where one theory eventually succeeds over the other as a theory for all cancers, to "division of labor," where each theory comes to dominate over different kinds of cancer.) The two theories are additive if they are compatible with each other but don't represent overlapping phenomena, so we can simply string together the theories and their implications. Or they could be complementary, engaging in interaction by virtue of representing shared phenomena. I won't settle this matter definitively here. (Indeed, it remains essentially unsettled in cancer biology.) But it is noteworthy that several proponents of SMT as well as CSC experts have tried to integrate the SMT understanding of cancer with new insights from CSC research (Hanahan and Weinberg 2011; Kreso and Dick 2014).[23]

[23] One way this integration might be achieved is if stemness is considered one of the hallmarks of cancer. Stemness may overlap with some of the hallmarks of cancer first emphasized by Hanahan

I've argued that SMT, as a theory of disease, is a unifying theory in the sense that it unifies (through abstraction) models of the pathogenesis of cancers. However, this does not imply that it provides a complete perspective on cancer because it does not unify all models of cancer. It unifies only those models visualizing particular spatial and temporal scales that represent the transformation of somatic cells into cancer cells. My account of SMT thus leaves open the possibility (which, given that there are other stages and aspects of cancer progression not represented by SMT, is a highly likely one) that other distinct theories or models provide a complementary partial perspective on the pathogenesis of cancer (Plutynski 2020).

In conclusion, SMT is a theory of cancer that views diverse malignancies collectively as diseases of uncontrolled cell proliferation, genetic alteration, and molecular dysregulation. In the late 1900s, it integrated research findings from different fields and on diverse etiologies for cancer, and unified models of the pathogenesis of various kinds of cancer. Since then, further attempts have been made to extend its principles to new cancer hallmarks. However, while some philosophers and historians have emphasized the historical and contemporary importance of theories of cancer and contagion, others have argued that the real revolutionaries are scientific frameworks, a counterproposal that I will now consider.

## 2.4. Theories Versus Frameworks

While theories of disease might seem to be imbued with vast and potent explanatory power, explanations for particular diseases might also seem to screen off general theories, rendering theories less relevant. Perhaps the more interesting structures for understanding progress in medical science are frameworks—disciplinary matrixes or research programs. In this section, I will examine this suggestion.

Thus far we have considered what makes for a theory of disease. But what purpose does a theory serve in infectious disease research or cancer research? What function is served by integrating and unifying research on a

and Weinberg (2000). For instance, the ability to self-renew indefinitely, a component of stemness, requires limitless replicative potential, one of their original six cancer hallmarks. It may also overlap with Hanahan's (2022) proposed new hallmark "phenotype plasticity and disrupted differentiation," in which differentiated cells can de-differentiate to become progenitor cells or in which progenitor cells can resist differentiation.

family of diseases? One obvious suggestion is that theories of disease explain a vast range of phenomena. Indeed, theories of disease can provide answers to questions about a family of diseases as a whole, such as the question of why cancer became more common in the new modern medicine. We saw in section 2.3.1 that multistage theory can provide a partial explanation because it represents the development of adult cancers generally. In contrast, a scientific model for a particular type of cancer is unable to answer this question about cancers as a family because it lacks the requisite scope.

However, setting aside the occasional question that demands a theory with wide scope, a family of diseases like the cancers is a heterogeneous collection of pathologies, and our medical interests are diverse, typically driven not by a theory but rather by specific problems or puzzles. It is plausible, then, that theories are not where most of the action is in investigating cancer (Plutynski 2018). Scientists and doctors often ask questions that could be answered by a theory but where knowledge about a particular disease provides a more detailed and desirable answer. As an illustration, in answering the question Why does the patient have a fever?, it is often more useful for a doctor to say that their fever is due to infection with the malaria parasite or due to colorectal cancer than to answer that it is due to "an infection with a pathogen" or "a cancer" because the former answers are more informative for patient counseling and disease management. Before more specific answers are available, a theory might serve as a second-best substitute, as when an unknown contagious syndrome or multiple tumors of unknown origin are explained by reference to an unknown germ or unknown primary tumor that has metastasized, respectively. For a doctor diagnosing an unknown contagious syndrome or a scientist researching one, germ theory provides a useful starting point for further investigation and for implementing generic infection control measures. But once a more detailed set of answers is available, it will at least partly take priority or "screen off" the more abstract theory. Research on diseases and their medical management is driven by a search for specific mechanisms or piecemeal causal knowledge (Plutynski 2018; Kincaid 2008).

It seems equally plausible that the ultimate goal when researching or diagnosing diseases is not to provide explanations but instead to suggest interventions. Cancer research, for example, is driven largely by medicine and society's therapeutic interests (Kincaid 2008; Laplane 2016). Around the turn of the twenty-first century, "translational research" and "translational medicine" were popular keywords for research efforts meant to translate

biomedical discoveries into medical interventions (Solomon 2015). Rather than highly general theories, developing new treatments for diseases presumably requires detailed knowledge of causal mechanisms and the ways we can intervene in them. Therefore the therapeutic aims of medicine might also cast doubt on the importance of theories of disease in medicine and, in turn, on their importance for understanding medicine.

Some historians, philosophers, and scientists have discussed "paradigms," "research programs," or "explanatory frameworks" for disease rather than (or in addition to) theories of disease.[24] For instance, Morange (1993, 1997) refers to the "oncogene paradigm," Plutynski (2018) describes the "mechanistic research program" on cancer, and Alessandro Blasimme et al. (2013) discuss the oncogene and CSC "explanatory frameworks." I will refer to these kinds of entities as "scientific frameworks." Although there are important differences between Kuhn's (2012) disciplinary matrixes, Lakatos's (1976) research programs, and other frameworks proposed by philosophers, one thing that all of these proposals have in common is that they include elements other than theory. For example, Blasimme et al. write, "An explanatory framework—like in the case of the oncogene paradigm—is made of causal patterns, [mechanism] schemata . . . intuitions, hypotheses, evidential standards and various bits of evidence and data coming from different experimental settings and instrumental devices" (2013, 377). Minimally, these frameworks all include a blueprint or direction for future research. Thus a "germ framework" or "somatic mutation framework" might provide an alternative to theories for capturing the thematic focus of a body of scientific research (on germs, on somatic mutations). For instance, after the early 1980s, when a central role was accorded to mutations in cancer genes involved in molecular regulatory networks, it provided a pattern—or framework—for future researchers who searched unceasingly for further cancer genes and sought to fill in the details of the mechanisms in which those genes participated (Morange 1993, 1997).

However, even though theories may not be all there is to a scientific framework, they are often recognized by philosophers to be one core component (Kuhn 2012; Lakatos 1976). Accepting that there are frameworks for infectious disease research and cancer research does not require us to deny that

[24] For examples, see Schaffner (1980), Morange (1993, 1997), Bishop (1995), Sonnenschein and Soto (1998), van Helvoort (1999), Carter (2003), Gillies (2005), Frank (2007), Blasimme et al. (2013), R. Weinberg (2014), Goodwin (2015), Plutynski (2018). While not all of these authors use terms like "paradigm" or "research program" in a technical philosophical sense, many of them do.

there are theories of disease. In fact, if there are scientific frameworks for disease, theories can play an important role within them. For instance, a theory of disease can endow a framework with its blueprint by specifying a common etiological pattern for all contagious diseases or all cancers, which then suggests a concrete direction for researchers in search of more specific mechanisms. It is an open question whether we can identify scientific frameworks in infectious disease research or cancer research,[25] and one that I will not settle here. My main point is that one of the main roles accorded to scientific frameworks for disease—providing a direction to a field of research—is equally fulfilled by a theory of disease as an etiological blueprint (on its own or within a scientific framework).

Carter (2003) argues that the Germ, Dissemination, and Distinguishability hypotheses are components of a bacterial theory of disease that is itself but one theory in a wider Lakatosian research program he calls "the etiological standpoint," which is oriented toward the search for causes of diseases generally. However, rather than empirical claims, Carter argues that these hypotheses function as methodological rules within the program, what Lakatos (1976) called "heuristics." I won't adjudicate whether Carter has faithfully applied Lakatos's ideas to theories of disease or whether Lakatos's ideas are well suited to this area of science in the first place; I will simply examine Carter's contention that the hypotheses of germ theory are methodological principles. Carter lays down the Germ Hypothesis (his "Bacterial Hypothesis") like this: "[E]very disease is caused by bacteria" (2003, 87). Instead of a literal empirical generalization, Carter reads the Bacterial Hypothesis nonliterally as stipulating a rule: search for a bacterial cause of each disease. Carter seems to suggest that because they are universal laws, the bacterial theory's hypotheses cannot be confirmed empirically, and the Bacterial Hypothesis as literally construed is obviously false anyway because not every disease is caused by bacteria (as was well known, even in the late 1800s); thus we had better not read the theory's hypotheses as empirical claims.

[25] Elements commonly associated with scientific frameworks include metaphors, inference schemes, methods, and exemplars. These elements are certainly found alongside germ theory and SMT. Germ theory is often associated with botanical metaphors, Koch's postulates (an inference scheme), inoculation experiments (a method), and exemplars such as Pasteur's research on spontaneous generation or Koch's research on the tubercle bacillus. Meanwhile, SMT is often associated with metaphors like renegade troops or interrupted electrical circuits, viral and nonviral transfection experiments as well as tumorigenesis assays (methods), and exemplars such as Bishop and Varmus's discovery of *v-src* and *c-src*. However, whether or not these various pieces form a cohesive package—a framework—is a question that cannot be answered merely by noticing their mutual association.

In order to avoid tumbling down a rabbit hole that might land us at the problem of induction, I will set aside Carter's suggestion that universal generalizations cannot be confirmed. However, it is worth noting that within germ theory, the Germ Hypothesis is not written like a generalization about all diseases, at least not with the pen of the authors I quoted earlier (even if in some versions it *is* written as a generalization about all contagious diseases). This takes some steam out of Carter's reasons for denying that the Bacterial Hypothesis is an empirical claim. Moreover, while I have suggested that germ theory forms an etiological blueprint, this does not necessarily mean that it consists of methodological rules because empirical claims can have the exact same implication as Carter's rules: if we accept as an empirical fact (or even as a working assumption) that all contagious diseases are caused by microbes, it will lead us to search for microbial causes for contagious diseases. Finally, *even if* we accept that germ theory and SMT consist in rules rather than claims, we could (like Carter) still consider them theories if we adopt an instrumentalism about theories of disease in which these theories are tools rather than explanations. The upshot is that Carter's interpretation of the bacterial theory gives us no reason to deny that medicine has theories of disease.

Turning now to arguments against the relevance of theories of disease, I agree with others that capturing detailed and specific causal mechanisms is a central goal of medical science. However, theories play an equally important role here as background knowledge for interpreting the results of research into molecular mechanisms. For instance, SMT allowed scientists to reinterpret the results of tumor virus research, no longer viewing these studies as identifying specific viral causes of cancer but rather as identifying particular cellular oncogenes (Morange 1997). Moreover, even when detailed and specific mechanisms begin to materialize, there will often remain gaps in our causal story that must be filled in with more generic causes. When the causal role of various oncogenes and tumor suppressor genes was first established, the mechanisms linking them to cell proliferation were (and remain) incompletely understood. Gradually, it was learned that they played a role in various cellular processes, such as growth signaling and cell cycle progression. Without knowing specifically how, SMT filled in the story abstractly by hypothesizing that the mutations must somehow disrupt the normal cellular process, leading to dysregulated growth signaling or cell cycle progression. While in principle a complete mechanistic explanation might completely screen off SMT, it may not always be in our interests to seek

further mechanistic detail. Rather than a temporary placeholder, a theory of disease may serve as the permanent glue that holds various mechanistic details together.

Likewise, despite any skepticism about their therapeutic relevance, theories of disease are potentially useful in therapeutic discovery. If knowledge of mechanisms is valuable for suggesting targets of medical intervention, then by serving as a blueprint for mechanistic research a theory indirectly asserts its therapeutic relevance. A theory of disease can also inspire general therapeutic strategies. Hanahan and Weinberg (2011) argue that SMT suggests a strategy of "hallmark-targeting cancer drugs." While they do not elaborate in detail, I read them as proposing that because the hallmarks of cancer are required for cancer progression across a range of cancers, disrupting a cancer hallmark has the potential to treat a range of cancers. (Classical nonspecific chemotherapies that target rapidly proliferating cells might be understood as an example of this strategy.) Given its genetic and molecular focus, SMT is also aligned with a therapeutic strategy common in precision oncology in which targeted therapy that binds to specific molecular pathology is tailored to patients with a particular genetic signature (see Chapter 9). On the other hand, one of the main motivations for CSC research and the CSC model of therapy is to overcome the common problem of cancer relapse after conventional treatment by adopting a different therapeutic approach, one that targets CSCs rather than the bulk of the tumor (Laplane 2016).

In conclusion, a theory of disease can fulfill one of the roles accorded to frameworks for disease by framework proponents: it can provide a direction for etiological research (by serving as an etiological blueprint). A theory of disease can even do so within a scientific framework, and thus theory proponents and framework proponents need not be at odds. A theory of disease can also perform other important functions that justify its relevance: it can help us identify and tie together causal mechanisms, and it can suggest both general and (indirectly) more specific therapeutic strategies. Therefore there is no reason to reject theories as irrelevant for the story of modern medical science in favor of an alternative story starring frameworks. In the previous sections, I identified two products of scientific research—the germ theory and SMT—that are justly considered theories of disease because they explain the occurrence of disease (they have causal-theoretical content), they integrate a diverse range of scientific findings (they are theory-level), and they unify models of the pathogenesis of a family of disease (they are about

disease). I will leave as an open question whether there also exist scientific frameworks for disease with additional elements beyond the content of these theories.

## 2.5. Molecular Reductionism

One source of continuity between nineteenth- and twentieth-century laboratory science and contemporary molecular biology is their potential for reductionism. The pathogenesis and pathophysiology of contemporary diseases is explained through biological parts of the organism, including (and especially) molecular pathways. This lens can fail to capture important causal features at higher levels of organization. As we will now see, theories of cancer and the debates surrounding them provide an informative illustration of the problem, allowing us to diagnose it more accurately.

Since the second half of the 1900s, the biomedical sciences, including molecular biology, have dominated medical research (Morange 1998), which led Engel (1977, 130) to famously describe a "biomedicine" or "biomedical model" with molecular biology as its "basic scientific discipline." By the time Engel was writing, cancer research had hopped on the "molecular biological bandwagon" (Fujimura 1988). SMT exemplifies the molecular turn: two of its hypotheses are centered on genes (understood as sequences of DNA) and molecular pathways. In Chapter 1, I suggested that the biomedical model should be understood as describing an emphasis on biological variables in modern scientific medicine and its tendency to reduce the patient to their parts. Laboratory medicine a century ago followed this model just as epidemiological medicine does today. But while laboratory medicine reduced the body to physiological, biochemical, and bacterial parts, epidemiological medicine breaks it down further into genes and molecules.

In what sense(s) does medicine embrace "reductionism"? While Engel's article is unclear and even confused about the relevant philosophical commitments,[26] one inchoate idea about medical reductionism found in his writing can be described as "explanatory." Diseases and their pathophysiology or pathogenesis are often understood through explanations that dissect the human body into its constituent parts and make intelligible how

[26] Engel (1977) saddles the biomedical model with both Cartesian dualism and physicalism, which are incompatible metaphysical positions. Diane O'Leary (2020) provides a disambiguation and analysis of Engel's metaphysical ideas.

disease is produced and sustained by relations and interactions among those parts. Decomposition of the organism into its *molecular* parts is the defining feature of *molecular* reductionism.

Kenneth Schaffner (2006) analyzes explanatory reductions like the ones typically found in molecular biology and the biomedical sciences as follows. Rather than "sweeping reductions" (sometimes called "intertheoretic reductions"; Nagel 1961) in which one theory is subsumed under another, Schaffner argues that "creeping reductions" are found across the biological sciences. These are piecemeal reductions in which a single phenomenon (rather than an entire subject domain like biology) is explained by parts at "lower levels." Imagine a decomposition hierarchy in which a whole organism can be conceptually decomposed into organs, which can be decomposed into tissues (including tumors), which can be decomposed into cells (including cancer cells), which can be decomposed into molecules (including cancer genes). A reductive explanation (explanans) refers to microstructures that are components of the (at least partially) decomposable system displaying the macro-level phenomenon that is being explained (the explanandum). There are also "connectability assumptions" relating micro-descriptions to macro-descriptions (e.g., sequence of DNA base pairs = gene; Schaffner 2006). Finally, Schaffner argues that the micro-descriptions explain the phenomenon by providing a "causal mechanical explanation," what others have called a "mechanistic explanation" (Machamer et al. 2000). In a mechanistic explanation, microstructures interact with one another and are organized such that they are responsible for the phenomenon (Illari and Williamson 2012). The way the parts are arranged and interact makes it intelligible why the phenomenon occurs.

As an example of a reductive explanation, an organism's having metastatic cancer (the macro-phenomenon) can be explained by cancer cells from the primary tumor traveling to distant sites via blood or lymph (microstructures); or the growth of a tumor (the macro-phenomenon) can be explained by disruptions in molecular circuits leading to uncontrolled cell proliferation (microstructures). Schaffner notes that reductive explanations in the biomedical sciences are often multilevel, meaning that some micro-structures are parts of other micro-structures within the explanation. In the previous example, molecular circuits are components of cells, which are both found in the explanation. However, in reductive *molecular* explanations, molecular mechanisms are mainly or exclusively invoked to explain the phenomenon.

We can illustrate this idea of a reductive molecular explanation in more detail using Figure 2.2, displaying the integrated circuit of the cell. One hallmark of cancer I discussed earlier is insensitivity to antigrowth signals and unrestrained cell cycling. A cancer cell's capability of resisting external signals that normally put a break on cell cycling and cell replication can be given a reductive explanation by zooming in and isolating a pathway in the top right of Figure 2.2, beginning with the anti-growth factor TGFβ. To simplify a complex story, TGFβ is an extracellular signaling molecule that through various pathways prevents the phosphorylation of protein pRb (encoded by the tumor suppressor gene *RB*). Phosphorylation, the addition of phosphate groups, reversibly inactivates pRb, preventing pRb from sequestering E2F transcription factors that otherwise promote the expression of genes that drive the cell through the cell cycle toward DNA replication, mitosis, and cell division. When pRb's action is disrupted by mutations in *RB*, which occur in most human cancers, pRb irreversibly loses its function, so that E2Fs—sequestered nevermore—constitutively promote cell replication. Thus the cell becomes insensitive to antigrowth signals from molecules like TGFβ because the inhibitory pathways connecting them to cell replication are "broken" (Massagué 2008).

Several philosophers have charged SMT with molecular reductionism, especially genetic reductionism (e.g., Bertolaso 2016).[27] "Reductionism" is a term that is often used pejoratively, especially in medicine and in the humanities and social sciences. The supposed problem is that reductive explanations that include only biological variables miss out on complex interactions at multiple levels of organization in the pathophysiology of diseases (Engel 1977, 1981). We should distinguish two elements to this criticism. First, by decomposing systems into molecules, including

[27] While some philosophers argue that SMT adheres to reductionism as a metaphysical presupposition such as genetic determinism, the thesis that genotype determines phenotype (Marcum 2005), others suggest SMT adheres to reductionism as a heuristic or methodological strategy (Malaterre 2007). This debate is replicated by cancer biologists. Robert Weinberg has often reflected on the reductionism of cancer biology in the late 1900s, typically emphasizing the reductionist's *approach* or experimental *methods* (1988, 2014; Hanahan and Weinberg 2011). In contrast, cancer biologists Sonnenschein and Soto (1998, 2000) construe their opposition to SMT in metaphysical terms: against the reductionism of SMT, they argue that cancer is an emergent property of tissues (the property of tissue disorganization). One reason for ascribing some sort of metaphysical reductionist commitment to SMT seems to be that it helps make sense of its reductionist methods of molecular biology and preference for reductionist molecular explanations (Sonnenschein and Soto 2000; Marcum 2005). However, there are other reasons for why cancer biologists might privilege molecular assays and explanations. For instance, they might think that genes are a fruitful research target because genetic explanations for cancer have special virtues or because we have good methods for identifying the causal roles of cancer genes (Plutynski 2018).

genes, reductive explanations might fail to explain phenomena that arise at the system level. Not only does the complexity of the typically nonlinear interactions exceed the ability of simple models like Figure 2.2 to predict the evolution of the system as a whole, but complex systems might also be "integrative": the function of the parts might actually depend partly on the organization of the system as a whole (Bechtel and Richardson 2010). This criticism takes issue with the reductionism of "molecular reductionism."

Additionally, by decomposing organisms exclusively into parts described by molecular biology and especially genomic variables, molecular reductionists exclude behavioral, social, and environmental variables. Behavioral and social variables are excluded because they are not molecular variables. Environmental variables are sometimes also excluded when they do not refer to molecules or when the mechanisms through which environmental molecules act on the body are unknown. And of course, none of these three kinds of variables refers to genes or gene products. This criticism takes issue with the molecularization of "molecular reductionism." I will first evaluate the reductionist charge against SMT before briefly turning to the molecularization charge.[28]

Engel's (1977, 1981) own alternative to the biomedical model's reductionism was a biopsychosocial model in which biological, psychological, and social variables at multiple nested levels interact to produce the phenomena of disease. Engel drew on earlier "systems thinking" to develop his model (Bertalanffy 1968; P. Weiss 1977). Today systems biology is a burgeoning field of biology studying complex networks of interactions among biological entities. One set of tools it uses is network analyses, in which large-scale networks are modeled as a whole to explain and predict the behavior of parts of the network or the global behavior of the network (Green et al. 2018). Through methods such as graph theory and dynamical systems theory, systems biologists can analyze whole-network features to make predictions about causal interventions or system dynamics. While there is discussion about whether the results of network analysis qualify as mechanistic explanations (Braillard 2010; Gross 2011; Matthiessen 2017), importantly the network analyst's more global strategies depart from a simple decomposition of wholes and localization of function to parts (Green et al. 2018) that

[28] Reflecting on medical reductionism in the twenty-first century, Jeremy Greene and Joseph Loscalzo (2017) draw a similar distinction between what they call the "network medicine critique of reductionism" and a "social medicine critique of reductionism."

characterizes the sample mechanistic explanation I just discussed involving pRb. Thus network analyses may evade the anti-reductionist's challenge motivated by the complexity of disease.

Cancer is a paradigmatically complex disease (Plutynski 2021). It is characterized by highly interconnected molecular networks like the one reproduced in Figure 2.2 in addition to networks at other spatial scales, and these networks evolve over time as cancer progresses. Relying exclusively on reductive explanations featuring isolated pathways may be insufficient for capturing the complexity of cancer. Some cancer biologists have accordingly wielded network analysis to understand how cancer phenotype and progression depend on patterns of connectivity among genes, transcripts, and proteins at the system or network level, rather than (or in addition to) trying to predict the behavior of the system reductively, or from the bottom up (Huang 2011; Barillot et al. 2020). For example, global patterns of network connectivity can explain why certain genes are "driver genes" in cancer that drive cancer progression despite an otherwise robust gene network, as well as predict (at least in principle) how global patterns of gene expression in a tumor will evolve (Green et al. 2018). These approaches to cancer complexity may surpass (or complement) reductive explanations in helping us understand and predict carcinogenesis.

It is not obvious that SMT is opposed to a network analysis of cancer genes and proteins, which after all essentially requires analyzing Hanahan and Weinberg's (2000) integrated circuit of the cell using nonreductive strategies—studying network features rather than isolated pathways. SMT posits that genes encoding proteins involved in networks regulating cell proliferation and motility are causally relevant to cancer, but it does not compel us to analyze these networks in one particular way. It is rather the use of the results of SMT-guided research into the molecular biology of cancer for constructing bottom-up mechanistic explanations that is reductionist in the explanatory sense. In fact, SMT may be crucial for systems approaches to cancer because it provides needed background knowledge, positing the kinds of phenomena in need of explaining (uncontrolled cell proliferation, cancer hallmarks) and the kinds of causal networks relevant to their explanation (networks of genes and proteins regulating cellular functions such as cell cycling). The anti-reductionist critique of cancer research motivated by the complexity of cancer is better directed at certain lines of research that ruthlessly generate reductive explanations of a bottom-up variety—and not at SMT itself.

The anti-reductionist might still insist that even if it need not commit cancer biology to a tireless search for reductive explanations, SMT visualizes cancer as fundamentally a disorder of uncontrolled cell proliferation, genetic mutation, and molecular dysregulation, transforming how medicine and society understand cancer. Because it views the individual as a collection of genes and molecules, in some remaining sense SMT is reductionist, and it may lead us to ignore important causes of cancer that are not as easily understood in genetic or molecular terms, including certain social and environmental causes. This critique is less about providing overly simple reductive explanations in the face of complexity than about what I earlier called "molecularization." A focus on genes and their molecular products in cancer draws attention and resources away from environmental, behavioral, or social interventions for preventing or treating cancer that could prove more effective or just (Tabery 2023; Chin-Yee 2024).

In response, it could be argued that SMT provides only one partial lens for understanding cancer, and cancer research can be (and is) multifocal, viewing cancer through a variety of perspectives rather than being wedded to one (Plutynski 2018, 2020). The real culprit is the overemphasis on SMT to the exclusion of all other theories or perspectives on cancer. Again, the critique is recast as a problem with how SMT is being used rather than with SMT itself.

However, the charge of molecularization against SMT might be more appropriate than this pluralist, multiperspectival response allows once we remember that as a theory of disease, SMT's main function is to provide a blueprint for etiological research (rather than a particular explanatory perspective on cancer). By specifying a molecular pattern for carcinogenesis, SMT encourages a molecular biological agenda for cancer research. And while it is more plausible that SMT takes nothing away from other explanatory perspectives as one perspective on cancer among many, it is less plausible that SMT takes nothing away from other research programs qua etiological blueprint. For while the main limits to the number of explanatory perspectives that medicine can manage are mainly cognitive (e.g., limits on our attention or interest), the limits to the number of research programs in which medicine can invest are much more material, consisting in financial and human resources. Therefore my account of disease theories from section 2.4 justifies directing at least some of the criticism of molecularization in contemporary cancer research and cancer care at SMT (to the extent that molecularization is the problem that critics claim it is).

## 2.6. Conclusion

Scientific medicine did not displace theories of disease from medicine but instead replaced speculative, sweeping theories of all disease with scientific theories about families of disease. Modern theories of disease unify models of the pathogenesis of specific kinds of disease within a family of disease. The germ theory of nineteenth-century modern medicine unifies contagious diseases, while the multistage and somatic mutation theories of the new modern medicine unify cancers. Multistage theory, an epidemiological theory, explains cancer incidence by modeling it as a product of multiple rate-limited stages, eventually taken to be genetic mutations. SMT explains cancer in individuals as resulting from uncontrolled cell proliferation, genetic alteration, and molecular dysregulation. Disease theories provide a blueprint for etiological research, allowing them to serve a similar function as scientific frameworks. Finally, cancer care and research are commonly criticized for their molecular reductionism. While we should distinguish bottom-up reductive explanations from SMT per se and lay the blame for any failure to capture the complexity of cancer mainly on the former, the SMT research blueprint *is* partly responsible for directing medical capital toward research into cancer's genetic and molecular causes (rather than, for instance, its social and environmental causes).

It has been argued that the influence of the nineteenth-century germ theorists extends beyond their theory and includes the construction of a general model of disease causation and classification, the "monocausal model." In Chapter 3, I examine this model and the way that epidemiology—by advancing a multifactorial theory of disease causation—has played a central role in restructuring modern medicine's disease concepts.

# 3

# Causation and Classification of Disease

## Summary

Infectious diseases are often said to have a specific necessary etiology—*the* cause of the disease. Meanwhile, noncommunicable diseases are often said to be multifactorial in their etiology. It has been claimed that this difference occurs because of the way infections are classified, using a "monocausal model," and that this classification bears theoretical and practical fruit. I argue that the monocausal model is a myth. Specific pathology rather than specific etiology was the key conceptual innovation of modern medicine and holds promise for explaining and intervening in disease. Meanwhile, a multifactorial theory of disease etiology was given to medicine by contemporary epidemiology and captures the idea that etiologic factors are generally neither necessary nor sufficient for their diseases.

It has been argued that establishing a universal etiology for infectious diseases helped give rise to modern medicine. This revolution is sometimes thought to involve the development of a new model of disease classification, the monocausal model, in which a particular kind of disease is defined according to a specific causal etiology, rendering that cause necessary for the disease's occurrence. In contrast, noncommunicable diseases adhere to a multifactorial model. However, rather than a multifactorial model of classification, epidemiologists in the 1900s developed a theory of multifactorial etiology, now known in its mature form as the sufficient-component cause model, that represents the way etiologic factors interact to cause all diseases. Moreover, as several problem cases illustrate, the monocausal model is not our current model of infectious disease classification. Nor does the monocausal model yield interesting general explanations for diseases or facilitate more effective medical interventions or research. Instead, a constitutive model represents the classification of communicable and noncommunicable diseases. The constitutive model characterizes a kind of disease according to the bodily condition of which it consists. The constitutive model can often create categories of diseases that have a universal etiology by defining

*The New Modern Medicine*. Jonathan Fuller, Oxford University Press. © Oxford University Press 2025.
DOI: 10.1093/9780190066178.003.0004

the disease according to a specific pathology (bodily disease structures or processes), and it does so for many infections and cancers. Specific pathology is often the right ingredient for making better medical explanations, medical research, and medical interventions.

## 3.1. From Multicausality to Monocausality and Back Again

Multifactorial thinking pervades contemporary epidemiology and medicine, from the way we describe modern diseases as having multiple and variable etiology (MacMahon et al. 1960; Krieger 1994) to the way we measure causal risk factors for diseases and customize medical classification, prognosis, and prevention based on those risk factors (World Health Organization [WHO] 2005, 2014). Nancy Krieger argues that "notions of multiple causation and multivariate analysis are so commonplace and so embedded in modern epidemiologic reasoning that they hardly merit discussion as a model or as an approach to understanding disease" (1994, 891). As an example of multifactorial thinking, the major modifiable risk factors for cardiovascular disease, including stroke, are smoking, obesity, physical inactivity, dyslipidemia, hypertension, diet, and diabetes mellitus (Hennekens 2015). Individually causal risk factors are not sufficient for disease (not everyone who smokes has a stroke), nor are they necessary (not everyone who has a stroke smokes).

Not only cardiovascular events like stroke and heart attack are multifactorial but also chronic diseases like diabetes and dementia, injuries like bone fracture, and even chronic symptoms like back pain. Cancers may be the most noteworthy multifactorial diseases of all. They are sometimes caused by infectious agents (especially viruses), other times by chemical carcinogens, and while genetic and epigenetic changes are typically required for their development, multiple genes are causally implicated in each case of cancer, and different genes can be drivers in different cases of the same type of cancer. In short, the characteristic ailments of the new modern medicine—noncommunicable diseases—are characteristically multifactorial in their etiology. In the second half of the twentieth century, epidemiologists hunted down risk factors for noncommunicable diseases just as bacteriologists had hunted down germs in the second half of the previous century—but now with population studies rather than petri dishes.

The multifactorial etiology of contemporary ailments seems noteworthy only when set against this historical background. In the late nineteenth and early twentieth centuries, the paradigm medical maladies were infectious diseases, which are often described as having a single universal etiology. Particular infectious diseases are caused by a specific germ. The specific germ is even necessary for the disease—without variola virus, one cannot contract smallpox. The etiology of several diseases of hormone or vitamin deficiency, from cretinism to beriberi, was also discovered during this era. Each of these diseases is similarly understood to result from a specific, necessary cause (the absence of a specific hormone or vitamin).

Alex Broadbent (2009, 2013, 2014) calls this turn-of-the-twentieth-century understanding of diseases the "monocausal disease model" to emphasize the privileging of one particular cause. In contrast, the model of disease popular among epidemiologists and public health authorities beginning in the second half of the twentieth century is a "multifactorial model" that recognizes the contribution of multiple causal risk factors to the development of each type of disease. Broadbent argues that the monocausal model is more a model of classification than a model of discovery. Scientists define a specific disease as the disease produced by a specific cause.

The monocausal model is thought to have supplanted a model of disease classification in which types of disease were defined constitutively, in terms of the components that comprised them. In the late 1700s, these components were typically symptoms, but by the middle of the nineteenth century they were often pathological anatomical lesions (Foucault 1973). For the simple reason that different causes can give rise to the same symptoms and lesions, the causes of these diseases were not singular and universal but multiple and variant. From this perspective, it looks like epidemiology and medicine have been dragged back to a bygone era when disease etiology was multifarious, complex, and unwieldy.

According to Codell Carter (2003), the monocausal model, or what he calls "the etiological standpoint," came to characterize modern medicine beginning in the late nineteenth century. The story often goes that a monocausal understanding of infectious diseases allowed for breakthroughs in their treatment and prevention (A. Evans 1993; Carter 2003). Jeremy Greene and colleagues write, "Motivated by breakthroughs in cellular pathology, pathophysiology, and especially bacteriology, doctors increasingly came to see diseases as specific entities, each with its own specific causes, manifested as characteristic syndromes. This new model prompted doctors to seek

therapies tailored to the disease and not the patient" (2012, 1080).[1] When the etiology of a disease is singular and necessary, its prevention seems simple and logical: prevent the cause to prevent the disease. In comparison, multifactorial etiology is often associated with complexity (WHO 2005). A list of risk factors for a disease suggests no obvious preventive strategy. Further, Broadbent (2009, 2013) argues that contemporary multifactorial thinking is theoretically impoverished because unlike the monocausal model, which provides a unifying causal explanation for all cases of a particular type of disease, "multifactorialism" fails to furnish general explanations for diseases, instead offering an assortment of causes. Thus it is perhaps a challenge for the new modern medicine that many contemporary diseases are multifactorial and not monocausal.

In this chapter, I explore the relationship between disease causation and classification in medicine.[2] Rather than the concept of a "disease," which is a much-studied topic in philosophy of medicine, my interest is in kinds of disease. Disease kinds play a central role in organizing medical research and practice. Thus the concept of a kind of disease is as important as (if not more important than) the concept of a disease to the story of modern medicine. Moreover, because the history of disease causation and classification appears to track the history of modern medicine through the nineteenth and twentieth century—with monocausal infectious diseases and multifactorial noncommunicable diseases mapping onto the older modern medicine and the new modern medicine, respectively—we can potentially gain insight into the evolution of modern medicine by studying disease causation and classification. It turns out that while there is much to be learned about the history of modern medicine from studying disease classification, it is not a story of shifting models of disease classification.

[1] Greene and colleagues (2012) note that like other "therapeutic revolutions" this one was complex and that new therapies often took decades to arrive.

[2] Some caveats are needed upfront. The history and philosophy of psychiatric classification is a large topic, and one I will not be able to discuss here. For collections of papers on psychiatric nosology, see Kendler and Parnas (2013, 2017). I will also not discuss formal biomedical classification systems or ontologies like the International Classification of Diseases or the Systemized Nomenclature of Medicine. These systems often serve distinct purposes, including facilitating administrative recordkeeping, the organization of biomedical information, and statistical analysis, and are not often used by clinicians as guides to diagnosis. I will instead use disease categories as defined and presented in clinical textbooks and guidelines, which *are* more often used diagnostically (Hucklenbroich 2014). Finally, I will discuss only the classification of kinds of diseases and not higher levels of classification in which disease categories are nested (e.g., families of disease like cancers or infectious diseases) nor lower levels of classification nested within disease categories (e.g., disease subtypes like HER2-positive breast cancer). Disease kinds are often treated like the medical equivalent of biological species, as a level of classification privileged in medical diagnosis.

In Chapter 2, I examined several theories of disease, including germ theory and somatic mutation theory. These theories explain the causation and pathogenesis of a family of diseases. In section 3.3, I propose that in the second half of the twentieth century, epidemiologists developed a theory of multifactorial disease etiology. However, the multifactorial theory is not a theory of disease of the same ilk as the germ theory or somatic mutation theory. It is a metaphysical theory, consisting in a few metaphysical assumptions about the causation of all diseases, that is primarily meant to help with the resolution of epidemiological problems. Nonetheless, the multifactorial theory reaches out beyond epidemiology and influences the way that doctors think about disease and its prevention in the new modern medicine.

The monocausal model is a myth; specific pathology rather than specific etiology was the key conceptual innovation of modern medicine and holds promise for explaining and intervening in disease. In section 3.2, I reconstruct the monocausal model of disease classification and defend my reconstruction against rival accounts. In section 3.3, I examine multifactorial thinking in the twentieth century and argue that it represents a theory of disease etiology. In section 3.4, I introduce the constitutive model as the standard model for classifying noncommunicable diseases, and in section 3.5, I argue that the constitutive model rather than the monocausal model is used to classify infectious diseases as well. In section 3.6, I argue that the theoretical and practical benefits of a monocausal model have been oversold. Finally, in section 3.7, I propose that the constitutive model can often establish universal etiology by defining a disease according to its specific pathology, and that doing so can reap the benefits attributed to the monocausal model.

## 3.2. The Monocausal Model of Classification

The monocausal model (also called the "etiological standpoint" and the "doctrine of specific etiology") is often said to be the model according to which medicine defines types of infectious disease. Philosophers have offered differing renderings of the model, and I will begin by providing my own minimalist reconstruction of the monocausal model and comparing it with some alternatives.

In *The Rise of Causal Concepts of Disease*, Carter writes, "The etiological standpoint can be characterized by the belief that diseases are best controlled

and understood by means of causes and, in particular, by causes that are *natural* (that is, they depend on forces of nature as opposed to the wilful transgression of moral or social norms), *universal* (that is, the same cause is common to every instance of a given disease), and *necessary* (that is, a disease does not occur in the absence of its cause)" (2003, 1). The first criterion, the requirement that the cause of the disease is natural, immediately suggests a strategy for *discovering* a disease's etiology: empirical research, especially research in the natural sciences. Meanwhile, the third criterion, that the cause of the disease is necessary, suggests a principle for *defining* a disease category: one should define the disease according to the cause that was discovered. (The third criterion implies the second criterion of universality: if a certain cause is necessary for the disease, then that cause will always occur whenever the disease occurs.) Carter understands the "etiological standpoint" as comprised of multiple disease theories and as comprising the core of an etiological research program, but I won't be concerned with these finer details of his account here; I will consider only the central idea of establishing universal and necessary disease causation.

The etiologic standpoint is a taxonomic ideal that Carter and many other authors believe guides etiologic research and faithfully describes our paradigmatic infectious diseases. Epidemiologist Mervyn Susser (1973, 23) argues that nineteenth-century discoveries by Louis Pasteur and Robert Koch "led to the redefinition and reclassification of many disease entities [disease types] by the criterion of cause. . . . By current definition, tuberculosis is caused by the tubercle bacillus." Similarly, epidemiologist Kenneth Rothman claims, "Necessary causes are often identifiable as part of the definition of the effect. For example . . . infection with the tubercle bacillus is a necessary cause for tuberculosis" (1976, 588). And philosopher Caroline Whitbeck notes that after the success of the germ theory in the nineteenth century "the name of the disease came to reflect the type of entity thought to cause it, the so-called etiologic agent, and etiology soon came to be definitive (i.e., to be regarded as essential) for those diseases for which it was known" (1977, 622).

More recently, Alex Broadbent (2009, 2013, 2014) has referred to this principle of disease classification as the "monocausal model." He too emphasizes that "[t]he special status that the monocausal model offers to certain causes is not an empirical status, but a conceptual one. Certain causes *define* the disease in question" (2013, 156). According to Broadbent's reconstruction, the monocausal model places a necessity requirement on the defining cause.

Adopting this requirement, we can represent the monocausal model as follows:

*a* is case of disease *D only if* an *E* caused *a*.

In a case of infectious disease or poisoning, *E* refers to a specific etiologic agent (a specific germ or a specific toxin, respectively); in a disease of deficiency, it instead refers to the absence of a specific agent, such as a specific nutrient. The two key features of *E*—and thus of the monocausal model—are *specific causation* and *necessary causation*. *E* is a specific cause because it refers to one particular kind of causal agent; it cannot refer to a disjunction of several kinds of etiologic agent, or else the disease would not be *mono*causal. *E* is a necessary cause because *D* occurs only if *E* caused it.

We should further clarify the notions of specificity and necessity on which the monocausal model relies and distinguish them from related concepts that sometimes travel under the same terms or that are sometimes associated with infectious diseases, such as universal etiology, sufficient causes, and specific effects. By "specific," I mean that disease types covered by the monocausal model have one particular or *specific cause* (which doesn't imply that cases of that disease lack other causes that operate alongside the specific one). The particular cause is a *necessary cause* of that type of disease in that, without it, cases of that type of disease cannot occur. As I discuss below, this necessity is conceptual and arises because the disease is defined by that cause (rather than, for instance, the cause necessitating the effect due to some natural causal law). Typically, the *etiology* of a disease refers to some distal or upstream cause such as a genetic mutation, and often more narrowly to an exogenous or environmental cause such as a pathogen, toxin, nutritional deficiency, or other environmental exposure. Thus, for the monocausal model, the "specific etiology" of a disease refers to a particular etiologic factor without which cases of that type of disease do not occur.

*Universal etiology* refers to an etiologic factor that causes all cases of a type of disease. Monocausal diseases have a universal etiology because a necessary cause of a disease is bound to be a universal cause of that disease. In section 3.7, we will see that it is also possible for a disease to have a universal etiology (or near enough) without having a necessary cause. It's worth noting that a necessary cause isn't necessarily a *sufficient cause*. When a cause is sufficient, it always results in the occurrence of disease. The monocausal model makes no commitment to the sufficiency of etiologic factors.

Exposure to a pathogen on its own is not sufficient for disease occurrence due to the possibility of host immunity preventing disease, for instance. Momentarily I will discuss a more restricted concept of "circumstantial sufficiency," where an etiologic factor is sufficient given certain other specified causes or background circumstances. Finally, "specificity" may also refer to a *specific effect* of a given cause. When a cause is specific in this distinct sense, it has only one effect. The monocausal model is similarly not committed to its etiologic factors having specific effects. Pathogens often have effects on their host other than causing disease; for instance, they can trigger an immune response.

Presenting the monocausal model in this form draws attention to its role as a model for *defining* particular disease types/taxa such as anthrax or typhoid fever, rather than a model of causal discovery. Applied to the example of anthrax, an instance of infection (*a*) is a case of anthrax (*D*) *only if* the germ *B. anthracis* (*E*) caused the infection. Although this necessity arises because *D* is defined in terms of *E*, we cannot define *D* in terms of just any factor. The factor we choose must be a cause of *D*.[3] Whether or not a particular factor is a cause of *D* is an empirical matter, to be settled through empirical research rather than by stipulation. The monocausal model is a form of causal selection that is applied after a cause of several cases of ill health is discovered. "Causal selection" is often discussed by philosophers in relation to a token effect when selecting the most salient cause—*the* cause—from among several causes of the effect (Hesslow 1988), as when the cause of a fire is determined to be the match-striking rather than the presence of oxygen. In contrast, the monocausal model operates at the level of *types* of effect: a cause is selected to be *the* cause of a type of disease, specifically by defining the disease according to the cause.

Although the condition that *E* caused *a* is necessary for *a* to be a case of *D*, it is not sufficient. *B. anthracis* can cause many things: an immune response in those who have been vaccinated against the bacterium, the death of livestock, public hysteria. These occurrences are not thereby cases of anthrax. As a model for defining diseases, the monocausal model as I have presented it is incomplete,

[3] When fitted to the monocausal model, disease kinds are what Muhammad Ali Khalidi (2021) calls "type etiological kinds" because all members of the kind share the same kind of cause. Khalidi argues that there are several reasons to employ etiological classifications, including enabling causal retrodictions and causal explanations. In section 3.6, I will discuss the relationship between the monocausal model and explanation. It is a distinct question whether diseases, etiologically (or otherwise) defined, are *natural* kinds. On diseases and natural kinds, see Lange (2007), Williams (2011), Khalidi (2013), and Plutynski (2018).

yet the constraint that it places on disease classification—the requirement of defining disease types according to a specific cause—is mighty, nonetheless.

To see how specific etiology depends on causal classification and not only on causal discovery, take Koch's 1882 announcement that the tubercle bacillus is the cause of tuberculosis regardless of what parts of the body are affected. Throughout the nineteenth century and earlier, various tuberculous conditions—that is, conditions associated with tubercles or small nodules—were recognized, including phthisis (what we now know as tuberculosis of the lung) and miliary (disseminated) tuberculosis. When Koch described his experiments on the causal role of the tubercle bacillus (now named *Mycobacterium tuberculosis*) in "The Etiology of Tuberculosis," he wrote of one disease encompassing various forms, including "phthisistic" and miliary. His experiments involved culturing the same bacterium from cases of phthisis and miliary tuberculosis and inoculating animals at various body sites, especially intra-abdominally, producing similar pathological effects each time, particularly the finding of tubercles with a characteristic appearance under the microscope (what we would now describe as "granulomatous"). Koch thus discovered that the same bacterium caused tuberculous infections of the lung as well as disseminated tuberculous infections.

However, Koch did not discover that tuberculosis is one disease due to one bacterium—that part he stipulated. By characterizing tuberculosis as an infection with the tubercle bacillus, all these conditions of different parts of the body were united as one disease. Koch could have instead continued to maintain that tuberculous infections of the chest and disseminated tuberculous infectious were different diseases, distinguished based on their anatomical location (phthisis and miliary tuberculosis, respectively). He could even have characterized tuberculosis as a disease of granulomatous lesions affecting any part of the body, but then he would have been simply wrong that tuberculosis is always caused by the tubercle bacillus because (as we now know) granulomatous lesions can have other causes, both infectious and noninfectious (Ohshimo et al. 2017).

In Chapter 2, I identified the germ theory of disease with three hypotheses, including a Distinguishability Hypothesis: a specific kind of disease is caused by the same specific kind of germ and no other ("same cause"), and a specific germ causes the same specific disease and no other ("same effect").[4]

[4] This "one cause–one effect" idea is one of two concepts of specificity articulated by James Woodward (2010). When a disease is maximally specific in this sense, it has one and only one cause

The first part of the hypothesis, "same cause," includes a claim about specific germ causation that is guaranteed to be true if we define a disease as the disease caused by a specific germ. The Distinguishability Hypothesis is not purely a classificatory principle—the monocausal model applied to contagious diseases—because it makes additional claims that go beyond the monocausal model, especially about the effects of a specific kind of germ (which, as we will see in section 3.5, are challenged by the example of HIV and tuberculosis). However, to the extent that "same cause" is true, it is true partly because of how contagious diseases are defined in modern medicine, as the example of tuberculosis shows.

Broadbent's reconstruction of the monocausal model places a second requirement on the defining cause, a circumstantial sufficiency requirement: "[G]iven certain circumstances, which are not sufficient to cause D, every occurrence of C causes a case of D" (2013, 150).[5] As Broadbent (2009) notes, in the absence of any further requirement beyond necessity it is possible that more than one cause could satisfy the monocausal model. He provides the example of oxygen: oxygen is needed for the survival of the patient and is thus necessary for the patient's having (any) disease *D*.[6] However, as a model for defining diseases, the monocausal model need not limit the number of potential causes that could in principle be used to define the disease; rather, it need define the disease only in terms of one of these potential causes, which it can do with only a necessity requirement. Moreover, there are reasons to doubt that our current model of infectious

of a certain kind (e.g., it is caused by one and only one kind of pathogen), and that cause has one and only one effect of a certain kind (e.g., it causes one and only one kind of disease). In Woodward's other, related concept of specificity, a disease is maximally specific when the function relating its occurrence to a specific cause is bijective.

[5] Broadbent (2009) offers a different formulation of the second requirement (discussed in Fuller (2018b)) that I will not consider here.

[6] It is not clear that a circumstantial sufficiency requirement succeeds in limiting the number of causes satisfying the monocausal model to one, at least in a nonartificial way. One factor likely to be found in the "given circumstances" listed in a circumstantial sufficiency requirement is "lack of a sufficient immune response." Then we would define a disease like tuberculosis as "the disease caused by *tubercle bacillus* given *lack of a sufficient immune response against the tubercle bacillus*," plus other circumstances. But now tuberculosis is defined in terms of at least two causes: the tubercle bacillus and lack of a sufficient immune response against the bacillus. This definition is logically equivalent to "the disease caused by *lack of a sufficient immune response against the tubercle bacillus* given *tubercle bacillus*. . . ." So the only way that a circumstantial sufficiency requirement limits the number of causes satisfying Broadbent's monocausal model to exactly one is by choosing one of these two necessary causes to stand apart from all other causes in the definition of the disease, and the monocausal model gives us no further criterion for selecting one over the other (as Broadbent 2009 points out).

disease classification includes a circumstantial sufficiency requirement (Fuller 2018b). Thus my formulation of the monocausal model (represented above) will satisfy only Broadbent's necessity requirement. In including only the requirement that the disease is caused by a specific necessary etiology, I capture a more minimal, widespread idea of the late nineteenth century, expressed by Carter (2003), Susser (1973), Rothman (1976), and Whitbeck (1977) above.

Lauren Ross (2018) discusses a principle similar to the monocausal model that she associates with nineteenth-century germ theorists, "the doctrine of specific etiology" (Dubos 1959). Ross argues that the doctrine includes two senses of causal specificity: "shared-cause specificity" and "single-cause specificity." Ross defines shared-cause specificity as follows: "[F]or all instances of disease D the same factor C or the same combination of factors ($C_1, C_2, \ldots C_n$) cause every instance of D in the sense of (I)" (2018, 12). Meanwhile, she defines single-cause specificity or "monocausal etiology" like this: "[F]or a given instance of disease D a single factor C causes D in the sense of (I)" (8). In both definitions, (I) refers to an interventionist account of causation: roughly, C causes D if and only if (in principle) manipulating only C would change D (Woodward 2003).[7] Ross suggests that infectious diseases adhere to these two senses of specific etiology.

Ross's concept of shared-cause specificity is that of universal etiology—that a certain disease always has certain cause(s)—and is implied by my minimal monocausal model. On the other hand, single-cause specificity is a distinct idea not found in the minimal monocausal model. It is likely not an idea held by all germ theorists, nor is it a very plausible one. For these reasons, I will not include single-cause specificity in my rendering of the monocausal model.

[7] Ross argues that the doctrine of specific etiology is best interpreted within an interventionist framework for causation. She rejects the association of the doctrine with claims about necessity or sufficiency, arguing that interventionism provides a better account of specific etiology than do necessity and sufficiency. However, interventionism need not be at odds with the necessary causation that I associate with specific etiology. Specific etiology as I understand it is about causal selection rather than the concept of causation. The two problems run in parallel. Whatever a cause is, the monocausal model makes a specific cause necessary for disease D by including it in the definition of the disease. I will remain neutral as to what it means for an etiological factor to be a cause of D. The only theory or concept of causation that would complicate my analysis of the monocausal model is one in which a cause is a necessary condition for its effect (because then the necessity that the monocausal model bestows upon a specific cause would be redundant).

First, single-cause specificity was likely not an idea held by all germ theorists. As Ross notes, earlier in the nineteenth century physicians distinguished between predisposing and exciting causes, where the former usually referred to intrinsic or constitutional factors such as heredity and the latter usually referred to external or environmental causes such as contagia (King 1991; Carter 2003). Multiple predisposing and exciting causes were thought to interact to produce a case of disease. This thinking is precisely the denial of single-cause specificity, and the germ theory did not replace it. Just as the germs of a plant require fertile soil to germinate, many germ theorists recognized that the germs of a disease require a "cooperative" host to cause infection (Tomes 1998). By the late 1800s, it was accepted that not everyone exposed to a germ developed disease and that vaccination was one factor that could often prevent disease; in other words, contagious germs were not the only factor causally relevant to the development of a case of contagious disease. A new generation of scientists sought to understand the constitutional factors that interacted with the germ in manifesting illness (Krieger 2011; Parascandola 2011b).

Second, infectious diseases simply do not obey single-cause specificity—germ theorists were right in thinking that host factors and vaccination are also relevant causes. Ross denies that host factors like genetics and immunity provide interventionist causal control over infectious diseases in the sense of controlling the presence/absence of the disease. However, as illustrated by real-world examples of latent infections and carrier states that develop into disease, several factors—including single genes in severe combined immunodeficiency disease, immune-suppressing drugs, and HIV infection—can control, in the interventionist's sense, whether one develops an infectious disease or not.

Ross also argues that vaccination targets the bacterial cause of a disease like anthrax, and thus, while it exerts causal control over the infectious disease, it does so only by exerting control over the bacterium. Ross's single-cause specificity allows that there are causal intermediates; it merely disallows factors off the causal chain linking the identified single cause with the disease. However, one way to describe the etiology of a contagious disease is to define Ross's variable C as exposure status with respect to the pathogen (unexposed = 0, exposed = 1) and variable D as disease status (non-diseased = 0, diseased = 1). Then vaccination status V (unvaccinated = 0, vaccinated = 1) exerts causal control over D directly rather than acting through C: vaccination does not control one's exposure to the pathogen; it controls the development of disease independently of exposure. In other words, vaccination is a

distinct cause of the disease.[8] From this point on, when I use the term "specific etiology," I will be referring to the singular, shared, necessary etiology of a particular type of disease (the minimal monocausal model).

In conclusion, the monocausal model is a model of disease classification that characterizes a kind of disease as resulting necessarily from a specific kind of cause. It is natural to wonder whether there is a rival "multifactorial model" of disease classification; if not, how should we understand multifactorial thinking in contemporary epidemiology and medicine?

## 3.3. A Theory of Multifactorial Etiology

While multifactorial thinking is now commonplace in epidemiology and medicine, it has a recent history in the rise of noncommunicable disease epidemiology in the second half of the 1900s. Epidemiologists grappling with the causes of cancer and heart disease had an inchoate understanding of multifactorial causation that was developed into an influential theory of multifactorial etiology.

The prototypical diseases described as adhering to the monocausal model are infections. Over the twentieth century, as noncommunicable diseases began to command more of the scientist's attention, the doctrine of specific etiology was gradually challenged. While inoculation experiments could isolate a single pathogenic cause of an infectious disease—a specific species of bacterium purified through serial culturing and visualized under the microscope—the etiology of cancer and heart disease could not be deciphered in this way. It was epidemiologists who would launch key studies on cancer and heart disease[9] and would come to understand

[8] To put it in the language of causal graphs, Ross implies that the relevant graph is a chain, V→C→D, which she stipulates does not violate single-cause specificity. However, as I just described the etiology of an infectious disease, the relevant graph is a collider, C→D←V, which *does* violate her condition of single-cause specificity. Elsewhere, she suggests a different sense in which a disease could be "monocausal": there could be one cause (and only one cause) that exerts a high degree of "probable control" over the disease's occurrence (L. Ross 2023). In other words, absent the cause, it's highly probable the disease won't occur, and given the cause, it's highly probable the disease will occur. Unfortunately, many infectious diseases are not "monocausal" in this sense either. For infections against which most people receive an effective vaccine (many childhood diseases) and for infections that are not particularly contagious, the disease probably will *not* occur given exposure to the pathogen.

[9] I do not mean to overlook the important role that the laboratory played in cancer research (discussed in Chapter 2) and in research on heart disease (D. Jones 2013). Moreover, there is a sense in which cholesterol played the role of a "germ equivalent" in research on heart disease from the middle of the 1900s: cholesterol was singled out as a central cause (if not *the* cause) of heart disease, and by the end of the 1900s medicine even had a "magic bullet" (statins) that could target cholesterol levels with precision (Greene 2008).

noncommunicable diseases not in terms of a single cause but instead in the language of multifactorial etiology. If the monocausal model was born in the laboratory, multifactorial thinking was nurtured by the epidemiological cohort study.

In the middle of the twentieth century, a multifactorial way of thinking about the etiology of infections was even starting to supplement a monocausal way of thinking about them among some epidemiologists. In a 1948 editorial in the *American Journal of Public Health and the Nation's Health*, Winslow wrote, "We must never be blinded by the over-simplified conception that we may discover the cause of any disease. The tuberculosis bacillus is not *the* cause of tuberculosis. It is *a* cause of that disease. You cannot have tuberculosis without the *Mycobacterium*; but half the people whose tissues have been invaded by the *Mycobacterium*, as shown by tuberculin tests, suffer from no disease in a clinical form" (854). Winslow noted that host factors, climatic/seasonal factors, and nutritional factors were other important causes, even if the mycobacterium is necessary for tuberculosis. In the 1950s, research on smoking and lung cancer as well as on coronary risk factors and heart disease further contributed to a "new modern" way of thinking about causation in epidemiology.

Richard Doll and Bradford Hill are two epidemiologists often credited with the research establishing that smoking causes lung cancer. They discovered that hospital patients who died from lung cancer were much more likely to have smoked cigarettes, and doctors who smoke cigarettes are much more likely to die from lung cancer (Doll and Hill 1950, 1954). However, not all of those with lung cancer smoked cigarettes, and not all of those who smoked cigarettes developed lung cancer. For some detractors (e.g., Yerushalmy and Palmer 1959), this suggested that cigarette smoke was not a true cause of lung cancer because cigarette smoke was neither necessary nor sufficient for its development and thus wasn't a "specific" cause.

But other scientists drew the conclusion that it was the idea of a necessary, sufficient, specific cause that needed to be abandoned rather than the hypothesis that cigarette smoke causes lung cancer (e.g., Hammond 1955). The latter camp instead adopted a multifactorial theory of etiology (Parascandola 2011b). While causal sufficiency goes beyond the tenets of the monocausal model as I understand it, multifactorial thinking tends to deny both necessity and sufficiency by positing that no etiologic factor causes every case of the disease and that no etiologic factor causes the disease without the help of other factors.

In a classic article detailing his now famous nine "viewpoints" for inferring causation from association, which were an expansion of the criteria used by the influential 1964 U.S. Surgeon General's Report on smoking and lung cancer, Hill writes, "We must also keep in mind that diseases may have more than one cause. It has always been possible to acquire a cancer of the scrotum without sweeping chimneys or taking to mule-spinning in Lancashire. One-to-one relationships are not frequent. Indeed I believe that multi-causation is generally more likely than single causation" (1965, 9). By noting that diseases like cancers can occur in the absence of single, specific causes, Hill diminishes the importance of necessary and specific causation in contemporary epidemiology.[10]

Reflecting on his and Hill's research on cigarettes fifty years later, Doll (2002) distinguished between *a* cause and *the* cause of a disease. He suggested that *the* cause was an appropriate expression only in reference to agents that are necessary and are either sufficient or satisfy Koch's postulates (e.g., the tubercle bacillus as *the* cause of tuberculosis). Doll recalled how the observation that lung cancer sometimes occurred in nonsmokers had been invoked by critics as counterevidence against the conclusion that smoking causes lung cancer, and he argued that "the critics had confused *the* cause with *a* cause and had not appreciated what is now obvious: that many diseases have several causes, in the sense that we now give the word—that is, agents that are associated with a disease in such a way that increased exposure to them is followed by an increased risk of the disease and decreased exposure by a decreased risk" (500). The critics had confused *the* cause (specific, necessary etiology) with *a* cause (multifactorial etiology).

Doll is ambiguous as to whether factors like smoking "increased risk" at the individual level or merely at the population level. When a cause is responsible for an increased risk of disease in a population rather than producing the disease in every individual in that population, it could be because its effect on individuals is governed by chance, or because it is insufficient for the disease in individuals and the other causes it needs to produce the disease vary among individuals in the population (or because the cause is *both* "chancy" and insufficient). Twentieth-century epidemiologists differed in

[10] Both the 1964 Surgeon General's Report and Hill's viewpoints include specificity of the association as one consideration. But rather than specific and necessary causation (the monocausal model), specificity here refers to a potential cause being associated with only one potential disease (which, anyhow, was violated in the case of smoking and was treated in both reports as merely one consideration among many rather than a sine qua non).

whether they attributed "probabilistic causation" at the population level to indeterminism or chance at the individual level (Parascandola 2011a).

Multifactorial thinking was also nurtured by developments in cardiovascular disease epidemiology, most notably the Framingham Heart Study, started in 1947. The Framingham Study was a prospective cohort study that recruited healthy participants in Framingham, Massachusetts. Initially focused on the incidence of coronary heart disease, it evolved to examine what it eventually called "factors of risk" or "risk factors" for heart disease, first landing on high blood pressure, high blood cholesterol, and left ventricular hypertrophy (Kannel et al. 1961). Cardiovascular or coronary risk factors are individual factors that are statistically predictive of the development of heart disease, heart attacks, and strokes (Aronowitz 1998). They were identified by following individuals over time and correlating the occurrence of cardiovascular events with these clinical factors. Framingham may not have introduced the term or concept of a "risk factor" into medicine (D. Jones and Oppenheimer 2017); indeed, the concept of a factor that raises the risk of disease in a population was discussed in the debates about smoking and lung cancer in the 1950s and 1960s. However, Framingham is widely credited with popularizing risk-related concepts and styles of reasoning in medicine (Aronowitz 2015). Chronic disease epidemiologists were soon engaged in a new "risk factor epidemiology" that some have considered to be characteristic of "modern epidemiology" as it emerged in the post–World War II era (Susser 1985; Parascandola 2011a; Morabia 2004; Broadbent 2013).

Part of Framingham's novel contribution was the demonstration that multiple risk factors together foretold a greater risk of heart disease than they did in isolation, a finding that was facilitated by new multivariate statistics. Aronowitz argues that the multivariate model "contributed to the widespread use of the term 'multifactorial' to describe the etiology of CHD [coronary heart disease] and other diseases" (2015, 87) and that soon medicine came to understand heart disease as due to the interaction of multiple risk factors. In contrast, Giroux (2012) argues that the main clinical aims of the study were to identify factors that marked those at higher risk of heart disease and to develop equations for predicting individual quantitative risk, rather than to identify etiological factors. (In Chapter 6, I examine this idea of a patient's individual risk.) However, regardless of the Framingham investigators' intent and work, Framingham risk factors were soon seen as loci on which to intervene to prevent heart disease through lifestyle interventions or drugs, signaling a new risk factor–based preventive medicine (Greene 2007;

Aronowitz 2015), and heart disease came to be seen as a model multifactorial disease. Once again, calling these causes "risk factors" is ambiguous about whether they are correlated with a risk of disease in a population because they are chancy causes or because other causes with which they interact vary in a population (or for both reasons).

As a slogan, multifactorial diseases are those that have several causes, often including interacting genetic and environmental factors. More specifically, each case of the disease is caused by the interaction of multiple causes, and different causes produce different instances of the disease, phenomena that Lauren Ross (2023) calls "multicausality" and "causal heterogeneity," respectively. High blood pressure and poor diet might interact to produce one case of stroke, while smoking and sedentary lifestyle might interact to produce another case. Ross (2018, 2023) suggests that not all diseases display multicausality and causal heterogeneity; for instance, infectious diseases typically do not.

However, diseases satisfying the monocausal model such as infections *are* multifactorial because while they may have a universal or necessary cause (such as a pathogen), that cause is insufficient for disease, and the other causes needed to bring about disease often vary among individuals. For example, as Winslow (1948) observed, many individuals who are exposed to the tubercle bacillus develop latent tuberculosis infection but not active tuberculosis disease. Progression from pathogen exposure to latent infection to active tuberculosis is likely due to a variety of viral virulence factors and host susceptibility factors that vary among cases (Pai et al. 2016).

The characteristic feature that distinguishes monocausal diseases from so-called multifactorial diseases like heart attack, stroke, osteoporosis, and prostate cancer is that so-called multifactorial diseases lack a necessary or universal cause. For example, while smoking is a cause of each of the above diseases, any of them can occur in the absence of smoking. But this is at least partly because medicine does not define multifactorial diseases according to one of their causes, which is a statement about our classificatory choices rather than the "natural" etiology of these diseases.

Broadbent (2009) refers to the negative characterization of multifactorial diseases given above as "bare multifactorialism" and suggests that modern epidemiology seems to adopt it. However, the idea of multifactorial etiology as it developed in the second half of the 1900s within epidemiology had positive content that went beyond bare multifactorialism. For cancer epidemiologists and chronic disease epidemiologists, multifactorial

etiology was not the mere absence of monocausal classification but rather a theory of disease etiology.[11] Thus multifactorial etiology is not an alternative to the monocausal model, even though it is sometimes presented as such (Broadbent 2009; L. Ross 2018).

In 1960, MacMahon, Pugh, and Ipsen's pioneering textbook *Epidemiologic Methods* introduced the metaphor of the "web of causation," an attempt to refine the idea of multifactorial etiology. The web consists of multiple causal nodes interacting with one another via different strands, with disease at the center of the web. The contrast to a causal web is a simple causal chain that begins with a single cause and terminates with the disease (Krieger 1994; Parascandola 2011b; L. Ross 2018). In the 1970s, other textbooks incorporated multifactorial thinking (e.g., Susser 1973). Then Rothman (1976) introduced what would become known as the sufficient-component cause (SCC) model, or "causal pie model," of causation in epidemiology, wielding the theory to solve several problems of epidemiological inference in his widely influential textbook, *Modern Epidemiology* (1986). The theory has since become a staple of the conceptual literature on causation in epidemiology.

According to the SSC model (Rothman 1976; Rothman and Greenland 2005; Rothman et al. 2008), what we normally take to be causes in epidemiology and medicine, such as smoking cigarettes or *M. tuberculosis*, are not sufficient for a health outcome but are rather components of "sufficient causes" (also referred to as "complete causal mechanisms"). Sufficient causes are sufficient because they inevitably produce the outcome. A sufficient cause contains multiple (typically, many) component causes, each representing either the presence or the absence of an etiological factor. (When the absence of a factor is part of a sufficient cause, that factor can prevent the outcome when present.) Component causes "causally interact" with each other in producing the health outcome.

Together, all the components of a sufficient cause are "minimally sufficient" for the occurrence of an instance of the outcome, meaning that if any of them were missing, then that instance would not occur (unless a different complete sufficient cause were present). However, a given kind of sufficient cause is rarely (if ever) necessary for the occurrence of the outcome because there are other kinds of sufficient cause—consisting of unique collections of

[11] Broadbent (2013) suggests that contemporary multifactorial thinking is due to epidemiology, but he does not characterize multifactorialism as a theory of disease etiology with positive content.

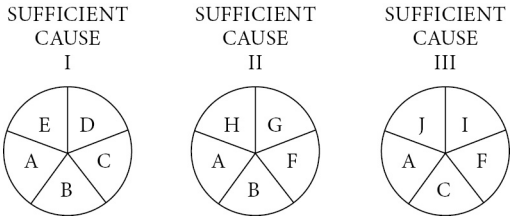



**Figure 3.1** The sufficient-component cause model.
Rothman (1976, 589).

component causes—that may be responsible for other instances of the outcome. A given kind of component cause is also not usually necessary for the occurrence of an outcome of that type. When a component cause *is* necessary (typically because it forms part of the definition for the outcome), it appears in every unique kind of sufficient cause (Rothman 1976). The underlying assumption is that a health outcome cannot occur in the absence of some sufficient cause, and if some component cause is an indispensable part of every sufficient cause, then that component is necessary for the outcome.

Rothman (1976) visualized the model using causal pie diagrams, which are an enduring image of multifactorial etiology in epidemiology and medicine (Figure 3.1).

Each pie represents one kind of sufficient cause of the outcome, while each pie slice represents one kind of component cause of the outcome.[12] In this example, slice A represents a necessary component cause. If the outcome is the disease COVID-19, then A would represent exposure to SARS-CoV-2, the virus that universally causes COVID-19. E might then represent lack of vaccination against COVID-19, while H might represent waning immunity from vaccination. Pies I and II include different sets of immune-related component causes that interact with viral exposure to cause COVID-19. In contrast, if the outcome was heart attack, there might be no environmental component cause A present in each and every causal pie.

In the third edition of *Modern Epidemiology* (2008), Rothman and coauthors framed the SSC model as a way of representing the "multifactorial etiology" of "all diseases" (15). They write that in the health and social sciences, "typical causes are neither necessary nor sufficient to bring

[12] Olaf Dammann (2020) argues in favor of regarding causal pies as complete etiological explanations and expands the SCC model to accommodate temporal dynamics.

about effects of interest. Hence, as has long been recognized in epidemiology, there is a need to develop a more refined conceptual model that can serve as a starting point in discussions of causation. In particular, such a model should address problems of multifactorial causation, confounding, interdependence of effects, direct and indirect effects, levels of causation, and systems or webs of causation. . . . This chapter describes one starting point, the sufficient-component cause model" (5). The SCC model is a theory of the multifactorial etiology of health outcomes primarily meant to solve theoretical, conceptual, and methodological problems faced by epidemiologists like the ones described by Rothman and colleagues. It differs from the theories of disease discussed in the previous chapter, such as germ theory and somatic mutation theory, which are each a theory about the etiology of a particular family of diseases. In contrast, the SCC theory represents the etiology of all diseases. Moreover, theories of disease like germ theory consist of empirical hypotheses or principles, those that can be confirmed empirically. On the other hand, SCC theory consists in metaphysical assumptions that are difficult to test empirically, like the assumption that health outcomes have sufficient causes and do not occur uncaused, where sufficient causes are sets of those factors that we normally call "causes." Thus it is best understood as a metaphysical theory about the etiology of health outcomes.

However, the SSC model is not a concept of causation or "theory of causation" as the latter expression is understood by some philosophers, in that the SCC model does not define what a cause is. It does not define a cause as a component of a sufficient condition for the effect. Confusion on this point is understandable because, as is recognized by epidemiologists (Parascandola and Weed 2001; Rothman et al. 2008), the model is similar to John Mackie's (1965, 1980) INUS account of causation, which *is* a theory about what a cause is. For Mackie, an INUS condition is "an *insufficient* but *necessary* part of a condition which is itself *unnecessary* but *sufficient* for the result" (1965, 245). A sufficient condition usually contains multiple INUS conditions. An INUS condition is "necessary" in that it is nonredundant within the sufficient condition, meaning that without it the sufficient condition is no longer sufficient (i.e., the sufficient condition is "minimally sufficient" in a sense near to Rothman's). Thus Rothman's component causes are INUS conditions, and his sufficient causes are minimally sufficient conditions.

However, Mackie (1965) *defines* a cause *A* of an effect *F* as "at least an INUS condition" for *F*.[13] In comparison, the SSC theory is not well understood as a definition of causation, although it is commonly interpreted as such (Parascandola and Weed 2001; Giroux 2015). In fact, Rothman and Greenland (2005) provide a separate definition of a cause.[14] SCC theory is better interpreted as a theory about the relationships among a health outcome and its causes.[15] The theory explains in broad strokes why there is heterogeneity in the occurrence of a health outcome within a population. For instance, Rothman uses the theory to show that the "strength" of an etiological factor (roughly, how much outcome it causes in a population) depends on the population's distribution of the other components of the sufficient causes in which the factor participates (its "cofactors"); also, the etiological fraction (the proportion of the outcome in the population caused by an etiological factor) for each etiological factor can (and usually will) sum to greater than 100% (Rothman 1976; Rothman and Greenland 2005). SSC theory can also be used to demystify certain mysteries of confounding and causal inference in epidemiology (Fuller 2019). And in Chapter 7, I will use it to help solve the problems of extrapolation from epidemiological studies. All these problems are explained by the way that etiological factors causally interact with one another according to SCC theory.

Moreover, the theory at least partly elucidates the probabilistic character of causation in epidemiology noted by epidemiologists like Doll (2002). When an etiologic factor raises or lowers the risk or probability of an outcome in a

[13] In other words, *A* is either an INUS condition for an effect that has multiple kinds of sufficient conditions, part of the sufficient condition for an effect that has only one kind of sufficient condition, or a condition that is itself sufficient for an effect. Mackie (1980) later adds that this definition is incomplete and needs an account of "causal priority," which he attempts to provide.

[14] "We can define a cause of a specific disease event as an antecedent event, condition, or characteristic that was necessary for the occurrence of the disease at the moment it occurred, given that other conditions are fixed. In other words, a cause of a disease event is an event, condition, or characteristic that preceded the disease event and without which the disease event either would not have occurred at all or would not have occurred until some later time" (Rothman and Greenland 2005, S144). While it is possible to interpret the first half of this underspecified definition using Mackie's theory of causation, it is also possible to interpret the second half using a counterfactual theory of causation. Elsewhere, Sander Greenland connects the SCC model to the counterfactual "potential outcomes framework" for causation in epidemiology (Greenland and Brumback 2002).

[15] Analogously, though Nancy Cartwright suggests that interventions used in evidence-based policy are INUS conditions or slices in a Rothman causal pie, she notes that hers is not a reductive account of causation because while causes are INUS conditions, not all INUS conditions are causes (Cartwright and Hardie 2012). Strevens (2007) and Wright (2011) develop nonreductive INUS accounts of causation.

population by less than 100%, this is because not all of the cofactors needed to complete a sufficient cause of the outcome are found in every member of the population. Therefore, at a population level, the risk or probability of the outcome will be altered by less than 100%, even if the individual risk of the outcome is altered by 100% in every individual in which those cofactors are all present (Rothman 1976). While the implicit causal determinism of this explanation for probabilistic causation at the population level is considered to be a drawback by those who believe that determinism is an unwarranted assumption in epidemiology (Parascandola and Weed 2001), the SCC theory can be modified so that sufficient causes are sufficient for a certain probability of the outcome in an individual rather than the outcome itself (Greenland and Brumback 2002).

If multifactorial etiology is a theory for all diseases, including infectious diseases, why did medicine develop it (or refine it) primarily in the context of epidemiological research on noncommunicable diseases in the middle of the 1900s? The web of causation and SCC theory are by no means the first theory of multifactorial etiology. We also encountered early nineteenth-century ideas about predisposing and exciting causes. But there was a partial interlude beginning in the late nineteenth century when contagious diseases came to be understood in terms of specific causes. During this interlude, scientists often tried to fit diseases like cancer into this framework, searching for viral causes of cancer (see Chapter 2) and rejecting smoking as a cause of lung cancer because it was neither necessary nor sufficient. Again, multifactorial thinking never completely went away (the germs of a contagious disease need fertile soil to germinate), but it didn't flourish either. Why did noncommunicable disease epidemiology precipitate its revival?

One potential answer is that as a pure factual matter, noncommunicable diseases have a more complex etiology than contagious diseases; they have more causes, and their causes vary more among individuals. Another potential answer is that the identification of a specific universal etiology for many contagious diseases exerted such a powerful pull on medical thinking that multifactorial thinking was partially displaced for a time. These answers are plausible.

However, it is also plausible that the methods of bacteriology oriented nineteenth-century bacteriologists toward specific universal etiology, while the methods of noncommunicable disease epidemiology oriented twentieth-century epidemiologists toward multifactorial etiology. Nineteenth-century bacteriologists like Koch used animal inoculation experiments to isolate

a *specific* (microbial) cause of a contagious disease through techniques like culture and microscopy that could purify and visualize the microbe. Koch's postulates, which were later articulated as a set of rules for determining whether a particular microbe causes a disease using inoculation experiments, included the rule that the (microbial) cause must be found *universally* in every case of the disease (Evans 1976). On the other hand, mid-twentieth-century "risk factor epidemiologists" used observational epidemiological studies that did not control for every cause of the disease except one, but rather measured the association of multiple factors with the disease. In the Framingham Study, they also used multivariate statistics to measure how these factors statistically interacted (Aronowitz 2015). The multifactorial theory helped epidemiologists interpret the results of epidemiological studies, just as ideas of specific etiology had helped bacteriologists interpret the results of laboratory inoculation experiments.

Multifactorial thinking about disease etiology developed within noncommunicable disease epidemiology in response to concrete empirical problems (Does smoking cause lung cancer?) as well as methodological and conceptual problems (Can etiological fractions add to greater than 100%?). Yet it quickly came to influence scientific medicine by offering a generic and flexible way of thinking about the etiology of disease, especially the etiology of noncommunicable diseases like cancer and heart disease, that counteracted the emphasis on specific causation more common in thinking about communicable diseases. Lists of causal risk factors now populate articles on the etiology of disease, and noncommunicable disease prevention often involves targeting multiple modifiable risk factors (e.g., diet, exercise) rather than using a silver bullet to target a singular cause.

Many of the risk factors associated with noncommunicable diseases today are genetic, including single nucleotide polymorphisms (single DNA base pair variations) and genetic mutations. Cancer, a stock multifactorial disease, provides a good illustration. According to somatic mutation theory (discussed in Chapter 2), multiple genetic changes interact to cause cancer by enabling cancer hallmarks that cooperate to give rise to fully malignant cancer cells. These genetic causes also interact with other bodily causes, such as immune factors, as well as environmental causes, such as mutagens and pro-inflammatory agents. Genetic risk factors are by no means limited to cancer. They are also central to the current precision medicine movement. In Chapter 9, I discuss how genetic risk factors are used to stratify patient groups in precision medicine and how multiple genetic risk factors are used

to predict the risk of disease in a multifactorial fashion using polygenic risk scores.

In summary, multifactorial etiology is not the denial of the monocausal model nor its own model of disease classification but instead an epidemiological theory of disease etiology that developed over the second half of the 1900s until it reached its mature form in SCC theory. The central tenet of SCC theory is that what we ordinarily call causes are interacting components of sufficient causes and that distinct sufficient causes are responsible for distinct cases of a disease. SCC theory reveals little about how contemporary noncommunicable diseases are classified, a problem to which I now turn.

## 3.4. The Constitutive Model of Classification: Noncommunicable Diseases

How are noncommunicable diseases generally classified in medicine, and why do they seldom have a universal etiology? It turns out that the answers to these two questions are related. Noncommunicable diseases are classified through a "constitutive model" according to the bodily condition in which they consist, and most noncommunicable bodily conditions lack a universal etiology.

First, recall that noncommunicable diseases lack a universal etiology at least partly due to our classificatory choices. There is nothing stopping us from defining a subtype of coronary artery disease (CAD) caused by smoking, just as we define a subtype of lung disease caused by *M. tuberculosis*. It is also tempting to think that smoking is merely a causal risk factor for CAD, which distinguishes it from *M. tuberculosis* and tuberculosis. But this thought is mistaken because *M. tuberculosis* is equally a risk factor for tuberculosis. Just as smoking imparts a risk of CAD of less than 100% in the population, *M. tuberculosis* exposure imparts a risk of tuberculosis of less than 100%. According to the multifactorial theory, in both cases the etiologic factor is but one component cause for the health outcome, and the reason it imparts a risk of disease of less than 100% is at least partly because its cofactors are not universally distributed in the population. This does not stop scientists from applying the monocausal model to tuberculosis, nor should it stop them from applying the model to CAD.

Non-monocausal diseases can be made to fit the monocausal model if we are willing to shift the borders of disease categories. In redefining

tuberculosis as an infection caused by the tubercle bacillus, certain states previously considered to be cases of tuberculosis were (whether knowingly or not) excluded from the newly reconstituted tuberculosis category. We could similarly redefine CAD as a coronary artery stenosis caused by smoking (CAD*). Cigarette smoke would then become a necessary cause of CAD*, and preventing smoking would invariably prevent CAD*. Preventing smoking would do nothing to prevent a coronary artery stenosis that falls outside of the reconstituted CAD* category, but then again, preventing exposure to the tubercle bacillus will do nothing to prevent symptomatically and pathologically similar infections that fall outside of the reconstituted tuberculosis category.

Noncommunicable diseases are not quite as resistant to the monocausal model as we might have thought. In order to achieve a monocausal classification system for noncommunicable diseases, scientists could simply redefine these diseases according to specific causes. But they do not. The nonuniversal etiology of many noncommunicable diseases thus results at least partly from our classificatory choices. This raises the question *What model* do *we use to define types of noncommunicable disease?*

Recall that for their lack of a singular etiology, noncommunicable diseases are often labeled "multifactorial diseases." In the early-mid 1800s, diseases were similarly recognized as having multiple causes, none of which was universally responsible for the disease (Carter 2003). I mentioned earlier that the reason for their multifactorial etiology was that diseases were defined constitutively, in terms of their component symptoms or pathological anatomical structures, and that several unique constellations of causes can give rise to the same symptoms and anatomy. How can we explain modern-day multifactorial etiology?

The answer is that like nineteenth-century maladies, noncommunicable diseases are classified according to their constitution—the bodily condition of which the disease consists. For instance, chronic obstructive pulmonary disease (COPD) is often defined as "*a disease state* characterized by airflow limitation that is not fully reversible" (Reilly et al. 2015, 1700, my emphasis). Osteoporosis is defined as "*a bone density* that falls 2.5 standard deviations (SD) below the mean for young healthy adults of the same sex" (Lindsay and Cosman 2015, 2488, my emphasis). Meanwhile, heart failure is defined by the American College of Cardiology as "*a complex clinical syndrome*" consisting characteristically of dyspnea, fatigue, edema, and rales (Mann and Chakinala 2015, 1500, my emphasis). Not to neglect acute cardiovascular conditions,

myocardial infarction is defined as "*myocardial necrosis* in a clinical setting consistent with acute myocardial ischemia" (Antman and Loscalzo 2015, 1602, my emphasis). Further, "[a] stroke, or cerebrovascular accident, is defined as *an abrupt onset of a neurologic deficit* that is attributable to a focal vascular cause" (Smith et al. 2015, 2559, my emphasis).

Turning to the realm of cancers, "The World Health Organization (WHO) defines lung cancer as *tumors arising from the respiratory epithelium* (bronchi, bronchioles, and alveoli)" (Horn et al. 2015, 507, my emphasis). There are several subtypes of lung cancer, including small-cell carcinoma, which is defined as follows: "Small-cell carcinomas *consist of* small-cells with scant cytoplasm, ill-defined cell borders, finely granular nuclear chromatin, absent or inconspicuous nucleoli, and a high mitotic count" (Horn et al. 2015, 507, my emphasis). In contrast, "mesotheliomas are *primary tumors that arise from the mesothelial cells* that line the pleural cavities" (Light 2015, 1717, my emphasis).

These examples illustrate a typical contemporary pattern: noncommunicable conditions are primarily classified according to what the condition is (e.g., a state giving rise to irreversible airflow limitation, a low bone density, a complex cardiac syndrome, myocardial necrosis, abrupt onset neurological deficits, tumors arising from respiratory epithelium, tumors arising from mesothelial cells). I will call this commonsensical taxonomic principle the *constitutive disease model*.

For diseases like COPD, osteoporosis, and small-cell lung cancer, the constitutive criterion is both necessary and sufficient.[16] We can thus represent the constitutive model as follows:

*a* is a case of disease *D* if and only if *a* is a *C*.

Applied to the example of osteoporosis, a bodily state (*a*) is a case of osteoporosis (*D*) if and only if *a* is a bone density that falls 2.5 standard deviations below the demographic mean (*C*). The key characteristic of *C* is *necessary constitution*. *D* only ever occurs when *C* occurs. In small-cell lung cancer, *C*

[16] For other disease types or subtypes, the constitutive criterion is only necessary, and additional criteria are needed to define the disease. For instance, type 1 diabetes is defined as (i) a disorder of glucose metabolism that is (ii) produced by the autoimmune destruction of pancreatic beta cells (Powers 2015b). The definition of heart failure requires that the syndrome results from ventricular dysfunction (Mann and Chakinala 2015). And in stroke the neurological deficits must be due to some focal vascular cause (Smith et al. 2015). So it is not the case that etiology or pathogenesis is irrelevant to disease classification.

refers to a tumor derived from the respiratory epithelium of the lungs that has particular morphological characteristics. In COPD, *C* refers to an irreversible disposition toward airflow limitation. The constitutive model is not wedded to a principle of specificity analogous to the monocausal model; the constitution of a noncommunicable disease is not always specific, as we will see in the next chapter. For instance, COPD can have different constitutive bases: sometimes it depends on enlarged alveoli, other times on narrowed bronchioles. In section 3.7, I will argue that when the constitution or pathology of a disease is specific, it can have a theoretical and practical payoff.

The constitutive model raises the interesting question Of what do diseases typically consist? What kind of a thing are diseases? As cancers of the thorax like small-cell lung cancer and mesothelioma indicate, a cancer is typically understood to be a malignant tumor of a certain cell type, though we could just as easily define it as a state of having this tumor developing or persisting in the body.[17] In Chapter 4, I will examine the nature of chronic diseases and conclude that they are generally dispositions (and that their dispositional nature helps explain why they typically lack a universal etiology).

In summary, noncommunicable diseases are typically not classified using the monocausal model but rather using a constitutive model, which defines a type of disease according to the bodily condition of which it consists. From the discussion so far, one might infer that our classification models on their own determine which diseases have a universal etiology: infectious diseases have one because they are defined using the monocausal model, while noncommunicable diseases lack one because they are defined using the constitutive model. Moreover, it might seem that the monocausal and constitutive models nicely demarcate late nineteenth-century modern medicine from the new modern medicine of today: while the diseases characteristic of late nineteenth-century modern medicine (infectious diseases) were defined using the monocausal model, those diseases characteristic of the new modern medicine (noncommunicable diseases) are defined using the constitutive model.

However, I will now argue that the monocausal model does *not* in fact describe the classification of infectious diseases today; the constitutive model does. Thus the difference between those diseases with a universal etiology

[17] Marta Bertolaso and John Dupre (2018) propose a process ontology of cancer in which cancer is the disruption of the normal balance of processes of cell proliferation, differentiation, and apoptosis.

and those diseases without one cannot be explained by the application of different models of disease classification.

## 3.5. Monocausal Model Versus Constitutive Model: Infectious Diseases

Carter calls the "etiological standpoint" (the monocausal model) "a defining characteristic of modern western thinking about disease" (2003, 1). On the received view, the monocausal model came to define types of infectious diseases beginning in the late nineteenth century.[18] But if ever the monocausal model (as characterized in section 3.2) was descriptive of infectious nosology, it is no longer. As the following problem cases illustrate, the monocausal model does not classify or incorrectly classifies many important infections, while the constitutive model succeeds where the monocausal model stumbles. I conclude that types of infectious disease are instead defined using the constitutive model; more specifically, they are defined according to the infection of which the disease consists.

The first set of problem cases includes infections that are not always caused by the same unique pathogen, and thus are not classified according to the monocausal model because they violate the requirement of causal specificity. Certain infectious conditions or syndromes can be caused by a great many different germs. The common cold is caused by many different viruses, and pneumonia is caused by many different microbes. While the monocausal model is usually described as a model of *diseases* (and syndromes are not always considered to be diseases), we might have thought that it would be applied to infections generally, given the praises often sung on its behalf. Yet some of the most common infections around do not adhere to its structure.

Even tuberculosis, an infectious disease that supposedly exemplifies the monocausal model, does not have a universal microbial cause. Robert Koch understood tuberculosis as a disease caused by the tubercle bacillus (A. Evans 1993). We now recognize tuberculosis as a disease caused by the *Mycobacterium tuberculosis* complex (MTBC) (American Thoracic Society 2000; Horsburgh 2015; Pai et al. 2016). *Mycobacterium tuberculosis* is the species most commonly responsible for tuberculosis, especially in North

[18] Susser (1973), Whitbeck (1977), Stehbens (1992), A. Evans (1993), Carter (2003), Greene et al. (2012), Broadbent (2014), L. Ross (2018).

America, but MTBC also includes other species of *Mycobacterium*, including *M. africanum* and *M. bovis*. It would be too nonspecific to say that tuberculosis is caused by mycobacteria because certain mycobacteria cause other diseases instead. Rather, MTBC consists of a subset of mycobacteria. But how can we appropriately define this subset? If we define MTBC as the mycobacterium that causes tuberculosis, then defining tuberculosis as a disease caused by MTBC would be circular. If instead we defined MTBC as a disjunction of *M. tuberculosis*, *M. africanum*, and so on, then MTBC does not refer to a *specific* etiologic agent.[19] Tuberculosis is supposed to be our paradigm monocausal disease, yet it seems to violate the monocausal model's condition of specific causation. Tuberculosis is not monocausal after all.

The constitutive model can handle these problem cases because specific causation is not a constraint on the constitutive model. In medicine, all viral infections of the upper respiratory tract are defined as "the common cold," and all infections of the alveoli of the lungs are defined as "pneumonia." The common cold and pneumonia are thus defined according to the state of which they consist: an infection of a particular part of the airway. Similarly, "tuberculosis" is an active infection with any of the MTBC mycobacteria.

When I say that an infectious disease consists of an "infection," I am referring to the state of being infected rather than the process of becoming infected. This allows us to say that tuberculosis is caused by MTBC bacteria; it amounts to saying that tuberculosis, the state of being infected with MTBC bacteria, is caused by exposure to MTBC bacteria or transmission of MTBC bacteria or the process of becoming infected with MTBC bacteria. It is also worth noting that the state of being infected with a microbe is not merely the state of having the microbe in or on one's body because uninfected people are colonized by microbes all the time. There are asymptomatic carriers of MTBC (who have "latent TB"), as well as asymptomatic carriers of many other pathogenic microbes, including *Streptococcus pyogenes*, the bacterium responsible for strep throat. An infection with a germ is best understood as the state of having the germ manifesting in the body in a certain

[19] It is not always clear what constitutes a specific agent. Influenza can be caused by the influenza A virus, the influenza B virus, or the influenza C virus. Is "the flu virus," then, a specific cause? One plausible test of specificity is that we must be able to usefully define the etiologic agent without resorting to a logical disjunction for that agent to count as specific. For bacterial infections, another plausible test is that the agent must refer to a particular species of bacterium. At present, MTBC consists of multiple mycobacterial species, and the only way to describe MTBC mycobacteria seems to be disjunctively.

characteristic way, which distinguishes mere colonization with the germ (or being a carrier) from infection (or having the disease).

The monocausal model is also thought to apply well to certain deficiencies (Carter 2003; Broadbent 2009). For these disorders, we might consider the specific cause to be the lack of a particular agent—a hormone or a nutrient. The cause of hypothyroidism might be a lack of thyroid hormone, while the cause of hypocalcemia might be a lack of calcium. However, just like infections, we can also understand these deficiencies using the constitutive model: as constituting the disorder itself rather than its cause. On this understanding, the causes of hypothyroidism (thyroid hormone deficiency) and of hypocalcemia (calcium deficiency) are several. For instance, thyroid hormone deficiency is sometimes caused by a lack of dietary iodine, and sometimes by autoimmune destruction of the thyroid gland. Thus the monocausal model does not apply uniquely to deficiencies, which can also be defined using the constitutive model.

The monocausal model runs into further trouble in classifying opportunistic infections. In these cases, it *does* classify the infection according to a specific cause, but it does so incorrectly. Consider opportunistic infections in patients with HIV disease. Infections are opportunistic when they result from a compromised immune system. In patients with HIV disease, the immune system is compromised because HIV infects T cells, which normally coordinate the immune response. One relatively common opportunistic infection in patients with HIV disease is an active respiratory infection with *M. tuberculosis*. It seems reasonable to define "HIV disease" according to the monocausal model: as an infection caused by HIV. Doing so would correctly classify the T cell infection as "HIV disease." But what about the respiratory infection?

As the respiratory infection was caused by *M. tuberculosis*, we could classify it as tuberculosis. However, the respiratory infection was also caused by HIV, which weakened the immune system and facilitated the opportunistic infection. Thus we should also define the respiratory infection as "HIV disease" according to our monocausal definition. Because infections with *M. tuberculosis* are unequivocally *not* cases of HIV disease, HIV disease must not be defined as an infection caused by HIV. Whenever one infection weakens the immune system, leading to a second infection, applying the monocausal model risks classifying the second infection incorrectly as the same disease as the first infection. Unfortunately, HIV infection can cause a wide range of opportunistic infections, so this problem for the monocausal model is not

limited to *M. tuberculosis* infections. More generally, when two pathogens contribute to the occurrence of the same case of infectious disease, the monocausal model may have trouble determining which disease has occurred.

Another problem of misclassification, described by Benjamin Smart (2014), is the case of human papillomavirus (HPV). HPV types 6 and 11 can cause rare infections like laryngeal papillomatosis and recurrent respiratory papillomatosis (RRP). The virus more commonly causes anogenital warts. If we define either laryngeal papillomatosis or RRP as an infection caused by HPV (type 6 or 11), then anogenital warts are misclassified as one of these two rarer conditions. Similarly, in addition to causing strep throat, *Streptococcus pyogenes* sometimes causes a skin infection called erysipelas. Defining strep throat as an infection caused by *S. pyogenes* misclassifies erysipelas as strep throat. More generally, when a particular pathogen can cause more than one type of infectious disease, the monocausal model may again have difficulty sorting out which disease has occurred in a particular case.

When infectious diseases are defined constitutively, these problems of misclassification do not arise. On the constitutive model, HIV disease is roughly defined as an infection with HIV, while laryngeal papillomatosis is defined as a laryngeal infection with HPV type 6/7, and strep throat is defined as an oropharyngeal infection with *S. pyogenes*. Then an opportunistic *M. tuberculosis* infection that was caused by HIV immunosuppression is *not* HIV disease because it is not an HIV infection, while HPV infection of the anus or genitals is not laryngeal papillomatosis because it is not a laryngeal infection, and *S. pyogenes* infection of the skin is not strep throat because it is not an oropharyngeal infection.[20]

In summary, the monocausal model fails to classify or incorrectly classifies several important infections. In comparison, the constitutive model avoids these problems and classifies tricky cases consistent with how they are actually classified by physicians. It does so quite simply by defining the disease

[20] Broadbent's (2013) version of the monocausal model, which includes a circumstantial sufficiency requirement (introduced in section 3.2), fares better than my more minimal monocausal model when confronted with the problems of opportunistic infections and pathogens that cause multiple diseases. However, the circumstantial sufficiency requirement has its own problems, making it equally unlikely that Broadbent's monocausal model describes contemporary infectious disease classification (Fuller 2018b). Briefly, the problems are threefold: doctors diagnose infections on the basis of the pathogen alone without identifying other causes comprising a circumstantially sufficient condition; circumstantially sufficient conditions are not known in full detail, which would prohibit diagnosis if they were part of disease classification; and compiling a full set of circumstantially sufficient conditions would require first identifying the disease for which they are circumstantially sufficient, which must not therefore be defined by those conditions. These problems also beset Smart's (2014) proposal to use their "full cause" to classify diseases.

according to the infection of which it consists. I conclude that the constitutive model is the standard model of classification for contemporary diseases of all sorts—noncommunicable as well as communicable.

One set of conditions that I have neglected so far are conditions defined symptomatically. Some of these conditions, such as chronic fatigue syndrome or the metabolic syndrome, are generally considered "syndromes"—that is, constellations of signs, symptoms, and clinical features—rather than "diseases," and the constitutive model is a model of disease classification. That said, syndromes also adhere to a constitutive model because a syndrome consists of the signs and symptoms that are used to define it. Other symptomatically defined conditions, especially chronic conditions such as epilepsy and gastroesophageal reflux disease, *are* generally considered diseases. However, as we will see in the next chapter, these diseases are also defined constitutively as the disposition that gives rise to their characteristic signs and symptoms.

We might wonder whether the monocausal model as I have described it was ever descriptive of infectious disease classification, or whether the constitutive model reigned throughout the nineteenth century and into the twentieth. If indeed the constitutive model operated throughout the microbiological revolution, then the dramatic shift in disease classification was not a change from an earlier "symptomatic model" of distinguishing diseases based on their symptoms to a monocausal model of distinguishing diseases based on their causes. Instead, the shift was from the constitutive classification of diseases in terms of the observable (symptoms, and then anatomy) to the constitutive classification of diseases in terms of the theoretical and unobservable (microbial infections). What changed was not medicine's model of disease classification but its conception of what diseases are, from a concept of diseases as clusters of signs and symptoms to a concept of diseases as theoretical entities.

In early nineteenth-century medicine, a distinction was often made between remote causes (what we now call "etiological factors" or "etiology") and proximate causes (what we now call "pathology" or "pathophysiology"; Carter 2003; King 1991). Remote causes outside the body brought about proximate causes inside the body, while proximate causes brought about symptoms. And when diseases were characterized by their symptoms, it made good sense to speak of both remote and proximate causes as "causes of the disease." In the 1700s, William Cullen defined phthisis (pulmonary tuberculosis) as the coughing up of pus or purulent matter along with a "hectic

fever" that spiked in the morning and evening and was often accompanied by facial flushing (King 1982). Cullen believed that the lung tubercles often found on autopsy in patients who'd had phthisis were a proximate cause of the disease. Then in the first half of the 1800s, diseases were increasingly redefined according to the anatomical lesion that produced the readily observable symptoms of disease (Foucault 1973; Taylor 1979). René Laennec redefined tuberculosis as the destructive changes found in the lungs, including the tuberculous matter. Lung tubercles—what had been a proximate cause of tuberculosis-as-symptoms—became the essence of tuberculosis-as–pathological anatomy. However, anatomical changes were still often referred to as the proximate cause of the disease, a conceptual confusion that did not go unnoticed. Thomas Watson lamented that "[t]he proximate cause is nothing else than the actual disease itself—the actual condition of that part of the body, from which the whole train of morbid phenomena essentially flows" and that the term "proximate cause" is "puzzling" and "tends to give the study of disease a scholastic and repulsive aspect" (1858, 76).

Plausibly, after the golden age of bacteriology at the end of the 1800s, germ theorists continued in this tradition, redefining each kind of contagious disease as the disease constituted by a proximate cause of the observable symptoms (namely, infection with a specific germ) rather than as the disease caused by a specific remote cause, as the monocausal model would have it.[21] After Koch, tuberculosis became infection with the tubercle bacillus. The mistake of considering this proximate cause to be the cause of the disease (rather than the cause of the disease's symptoms) is then the same one that Watson criticized in 1858: the error of identifying a disease as its own cause. This confusion may continue today, perpetuating the myth of the monocausal model. For instance, *Harrison's Principles of Internal Medicine* (Kasper et al. 2015) defines influenza as "an acute respiratory illness caused by infection with influenza viruses" (1209) and cholera as "disease caused by *V. cholerae* serogroup O1 or O139" (1061).

[21] Carter agrees: "At the beginning of the nineteenth century, whenever possible, diseases (which had earlier been construed as collections of symptoms) were redefined in terms of specific lesions; with the rise of the etiological standpoint, whenever possible, diseases (then construed as lesions) were redefined in terms of bacterial infections. In both cases, the new definitions created proximate causes that were universal and necessary" (2003, 107). But then the etiological standpoint is misnamed and misdescribed because it is really concerned with proximate pathology rather than remote etiology. Moreover, this passage seems to contradict Carter's central claim that the etiological standpoint arose later in the century with bacteriology (rather than earlier in the century with pathological anatomy).

In short, it turns out that the monocausal model is not the model that doctors use to define types of infectious disease (and may never have been); rather, the constitutive model classifies types of infectious disease according to the infection of which the disease consists.

## 3.6. The Myth of the Monocausal Model

I have argued that the monocausal model does not guide the classification of infectious or noninfectious diseases. However, we might wonder whether the monocausal model is better seen as a useful ideal that medicine should implement (but currently does not), a prescriptive model rather than a descriptive one. The "monocausal model," "doctrine of specific etiology," or "etiological standpoint" is often seen as a great theoretical and therapeutic triumph of modern medicine. According to Carter, "[o]f the numerous changes that have occurred in medical thinking over the last two centuries, none have been more consequential" (2003, 1). In the middle of the 1900s, during a golden age of medical discovery, René Dubos wrote, "Unquestionably the doctrine of specific etiology has been the most constructive force in medical research for almost a century and the theoretical and practical achievements to which it has led constitute the bulk of modern medicine" (1959, 102), while adding that few diseases adhere to its strictures. More recently, F. Kräupl Taylor suggested that "[m]ost diseases which are acknowledged in medicine today are certainly not monogenic disease entities [those with a specific etiology or 'pathological onset']. Yet the final hope and aim of medical science is the establishment of monogenic disease entities" (1979, 21).

It might be thought that the monocausal model can correct for the theoretical and therapeutic shortcomings of the multifactorial theory by unifying disease explanations and interventions around a single causal factor. However, the theoretical and practical benefits of a monocausal model have been overstated. The monocausal model is a myth, both descriptively inaccurate and prescriptively underwhelming.

The theory of multifactorial etiology makes sense of several phenomena in epidemiology and medicine. For instance, it helps explain the variation in outcome of exposure to risk factors and medical interventions among members of a population. It highlights the role of causal interaction in producing this variation, as well as that multiple mechanisms can produce the same outcome. However, it fails to discriminate among the component

causes of disease and to pick out any specific causes from among the lot, leading to a complicated etiological perspective and no unifying explanation for the occurrence of that disease. Moreover, ceteris paribus, preventive strategies that target a nonuniversal component cause (such as lowering blood pressure to prevent a heart attack) might have a more limited effect on the population burden of the particular disease or health outcome, at least compared to strategies that target a universal etiology (such as vaccinating against a pathogen to prevent an infectious disease). From an individual person perspective, ceteris paribus, preventive strategies that target a nonuniversal etiology might similarly have a lower probability of preventing a particular disease compared to strategies that target a universal etiology. I will first consider the monocausal model's purported value in explaining disease.

Broadbent argues that the "central theoretical advantage of the monocausal model of disease, is that it encourages *general explanation*" (2009, 304). By "general explanation," he means that the monocausal model identifies a cause common to all cases of a disease type that thus provides an explanation for why all cases of that disease type occur, why the disease type occurs in *general*. He argues that a weakness of a multifactorial approach that does not define a disease type according to particular cause(s) is that it will often fail to achieve this theoretically and practically worthwhile general understanding because it does not foreground causes that can serve in a general causal explanation. In response, he proposes a contrastive model of disease that defines a disease type via a contrast between cases with certain defining symptoms (Broadbent 2009) or dysfunction (Broadbent 2014) due to a set of causes (those who by definition have that disease), and a contrast class of people without those symptoms/dysfunction or causes (those who by definition do not have the disease). Broadbent argues that his contrastive model shares with the monocausal model the virtue of providing general explanations while avoiding the vice of restricting the model to a single cause.

However, neither the monocausal model nor the contrastive model can provide satisfactory general explanations because they stipulate the generality that they seek to explain. A virtue of a general explanation is that it explains what otherwise might be surprising: why all members of a certain type share some feature. A satisfying explanation thus adds to our understanding of why those members all share that feature. Broadbent (2009) provides the example of the law of gravity, which explains why all earthly masses tend to fall to Earth through the synthetic fact that Earth's gravity acts

on all masses near its surface. However, the monocausal (and contrastive) model achieves generality by defining it into being, and in so doing provides an explanation that is neither surprising nor adds to our understanding. On the monocausal model, to explain that all cases of cholera occur because they are caused by *Vibrio cholerae* is to say that all cases of the disease caused by *V. cholerae* occur because they are caused by *V. cholerae*. This answer fails to provide an interesting scientific explanation for why all cases of cholera occur. The reason that all cases of cholera share this cause is because we've stipulated that they do, so we should not be surprised by the fact that all cases share this cause. Nor does the explanation add to what we already know about cholera because the causal information it provides about cases of cholera is already provided by cholera's definition.[22]

On further reflection, a monocausal model also fails to deliver many of its promised practical benefits. It is tempting to award credit to the monocausal model for medical breakthroughs like vaccines and antibiotics, which target the universal offending pathogen. Alfred Evans recounts the following anecdote. In 1882, Robert Koch presented convincing evidence to the Berlin Physiological Society that the tubercle bacillus he had isolated was the cause of tuberculosis. One of the many esteemed attendees of the meeting was Paul Ehrlich, who later opined, "That evening was engraved in my memory as the most majestic event I have ever participated in" (quoted in A. Evans 1993, 21–22). Ehrlich went on to lead a research laboratory that synthesized the world's first somewhat effective and specific chemotherapy for an infectious disease: arsphenamine (tradename Salvarsan), a treatment for syphilis. But once we separate the monocausal model from the germ theory, which certainly *did* play an important role in the discovery of treatments like arsphenamine, it is less clear that the monocausal model is a good guide to intervention or a good approach to organizing medical research.

First, the monocausal model is a suboptimal guide to intervention, including rational prevention. Caroline Whitbeck (1977), Codell Carter (2003), and Lauren Ross (2018) claim that classifying a disease according to a particular etiologic agent serves medicine's interest in preventing disease.

[22] On the other hand, the monocausal model can facilitate a token causal explanation for the occurrence of a case of disease. Classifying a token illness as a case of a particular kind of monocausal disease serves as a shorthand for a token causal explanation. For example, on the monocausal model, classifying a case of diarrheal illness as cholera would imply that the case was caused by *V. cholerae* (see Khalidi 2021 on etiological kinds and explanation). However, while the monocausal model can facilitate a token causal explanation for each case of a type of disease, it fails to provide a satisfactory general causal explanation for the disease type, for reasons I just articulated.

How so? "If we wish to prevent something, and we do generally wish to prevent diseases, we need only prevent or eliminate something which is a necessary condition for it," Whitbeck argues (1977, 629). By comparison, the multifactorial theory selects no cause to be necessary for the occurrence of a disease. We can call this principle, that preventing a necessary cause of a disease will necessarily prevent the disease, "modus tollens prevention" (if disease, then cause; no cause; therefore, no disease). Modus tollens prevention is a logically sound procedure. But logical necessity is not always a *strategically* sound guide to action. If your single male friend complained to you that they were tired of being a bachelor, it would be somewhat unhelpful to suggest to them that they are *guarantee*d to cure their bachelorhood if they just stop being an unmarried man (given that being an unmarried man is a necessary condition for being a bachelor). Similarly, the principle that we can prevent tuberculosis by preventing MTBC bacteria from causing tuberculosis or that we can prevent CAD* (coronary artery disease caused by smoking) by avoiding smoking is unhelpful.

The problem with modus tollens prevention is that the kind of causal necessity provided by the monocausal model—conceptual necessity—is not instructive for those wishing to prevent disease. We can grant that knowing about the cause of a disease is often the first step toward rationally preventing the disease. But there is little further value to be gained by knowing that this cause is necessary for the disease if the necessity simply results from definition. If we already know that MTBC bacteria cause infections or that smoking causes coronary artery stenosis, then we learn nothing new of practical relevance by being told that preventing MTBC bacteria will necessarily prevent tuberculosis (because infections caused by MTBC bacteria are known as tuberculosis) or that preventing smoking will necessarily prevent CAD* (because coronary artery stenosis caused by smoking is known as CAD*). We *can* learn a great deal of practical relevance through causal discovery, by learning that MTBC bacteria cause infections (for instance). But again, the monocausal model is a model of causal selection that kicks in after a cause has been discovered.

Second, the monocausal model is not always the best or even a very good approach to organizing medical research or therapy. It's tempting to think that targeting necessary causes will have a greater public health and medical impact than preventing non-necessary causes. Even if modus tollens prevention does not help us rationally devise preventive strategies, once we have found a way of eliminating or neutralizing a necessary cause, then deploying that

strategy is guaranteed to prevent all cases of the disease, which cannot occur without the necessary cause. Our preventive strategies will thus be more effective for monocausal diseases compared to diseases that lack a necessary etiology. However, this is a mistaken line of reasoning, again because of how the monocausal model obtains necessary causes. To see why, recall that preventing smoking will not prevent all cases of CAD. Making smoking necessary for CAD* through stipulation implies that preventing smoking will prevent all cases of CAD*. But doing so will not help us prevent more cases of the coronary artery stenosis that disposes people to the heart attacks that we ultimately wish to avoid, because CAD* simply excludes cases of coronary artery stenosis that are not caused by smoking—it does not make them magically disappear.

One might think that even if the monocausal model does not provide good guidance for rational prevention or improve the power of medical prevention, when research is organized around particular kinds of disease the monocausal model ensures that these are the right kinds for facilitating the development of specific preventive or curative interventions. While the monocausal model does not suggest that inoculating an individual with a specific attenuated pathogen or its antigens will provide immunity to a particular infectious disease, when that disease is defined according to the specific pathogen, vaccine research on the disease can properly yield a vaccine that is specific for the disease. We can then prevent the disease using the specific vaccine. If we don't distinguish between influenza and COVID-19 but study these diseases together under the category "pandemic diseases caused by respiratory viruses," then vaccine research will likely not yield a vaccine that is effective for this broader category of respiratory illness because a particular vaccine will typically generate immunity against an antigen specific to a particular pathogen. When we instead develop vaccines against COVID-19 and against influenza, where these diseases are defined as infections caused by SARS-CoV-2 and by influenza virus, respectively, we happily end up with specific and effective vaccines.

However, in other examples, especially examples of treatment rather than prevention, the monocausal model reveals itself to be of limited value because specific etiologic factors do not always cause specific bodily pathology. For example, because smoking causes an anatomically and histologically diverse range of cancers (from lung cancer to bladder cancer), using the monocausal model to define a category of "cancer caused by smoking" would yield a disease too motley to usefully organize research on cancer treatment. The problem is that even though smoking would be a specific cause of this family

of cancers, its action in the body would be nonspecific to cell and tissue type, anatomical location, and genetic mutation, yet effective treatment often depends on cell and tissue type, anatomical location, and genetic mutation. We will see in the next section that a key to unlocking specific treatment is a different kind of specificity: specific pathology. In fact, specific pathology can often achieve the universal etiology that has proven so useful in disease prevention.

The monocausal model is not only descriptively inadequate for understanding why many communicable diseases have a universal etiology and why many noncommunicable diseases lack one. It is also of limited theoretical and therapeutic value as a prescriptive model. We will see in the next section that any virtues of the monocausal model for prevention and treatment are often matched by the constitutive model when we define a disease according to a specific pathology.

## 3.7. Specific Pathology as a Guide to Universal Etiology

In sections 3.4 and 3.5, I argued that communicable and noncommunicable diseases are classified using the same model, the constitutive model. Despite popular philosophical and historical opinion, the monocausal model is not used to impart communicable diseases with a specific universal etiology by defining them according to a specific cause. In section 3.6, I argued that the monocausal model is not as lucrative as it is sometimes made out to be, anyway, because it creates necessary etiology through stipulation. Yet communicable diseases often *do* have a universal etiology. And their universal etiology is often theoretically and pragmatically valuable. In this section, I argue that diseases can often obtain a worthwhile universal etiology when the constitutive model classifies them using a different type of specificity: specific pathology.

While etiology refers to a cause of a disease (often a cause external to the body, the "remote cause"), pathology refers to disease structures or processes within the body (often those that constitute the disease or serve as the causal basis of the disease, the "proximate cause").[23] For instance, cancers are

[23] While a distinction is often made between pathology and pathophysiology, where "pathology" refers to structures and "pathophysiology" to processes, this distinction is not important here, where I use "pathology" as a more general term that covers both structures and processes. I recover the distinction between pathology and pathophysiology in Chapter 4.

typically defined according to a specific histopathology, the specific cell type and organ from which the tumor arises, a practice originating in the 1800s (Triolo 1964).[24] In the early nineteenth century, diseases of all sorts were increasingly classified according to their specific pathological anatomy, such as tubers of the lungs in tuberculosis. Thus the notion of a specific disease entity originated with pathological anatomy rather than later, with the germ theory (Rosenberg 2002), and was never about etiology so much as it was about pathology. The later nineteenth-century germ theorists simply extended the principle of specific pathology to infections by classifying types of infection according to the specific germ infecting the body. Koch stipulated that infection with the tubercle bacillus (anywhere in the body) was tuberculosis, while infection with the comma bacillus was cholera.

If infectious diseases are not classified using the monocausal model, and the principle to which they generally adhere is specific pathology rather than specific etiology, how do we explain why they *do* seem to have a universal etiology, as was noted by the germ theorists? It's first worth clarifying what the specific universal etiology of a contagious infection is, exactly. Recall that specific universal etiology refers to a particular upstream (usually) environmental cause of all cases of a kind of disease. Necessary causes are necessarily universal causes of their disease. But not all universal causes of a disease are necessary, resulting from the way we define the disease.

When we say that the cause of a particular disease is its specific pathogen, one plausible interpretation of this statement that avoids the conceptual error of identifying a disease with its own proximate cause goes like this: the remote cause of the disease is environmental exposure to the specific pathogen, or transmission of that pathogen to the host. It is indeed the case that many contagious diseases are universally caused by exposure to a specific germ in the environment. But the constitutive model does not necessitate this fact; it is not analytically true that infection with a specific germ is caused by environmental exposure to that specific germ. It could have turned out that infectious diseases arise through spontaneous generation of the pathogen in vivo, as many nineteenth-century thinkers thought; such a possibility is not ruled out when we define infectious diseases constitutively. In fact, certain

[24] Anya Plutynski (2018) argues that cancer classification is pluralistic: there are several ways of carving up cancer according not only to cell type and organ but also stage/grade and genetic subtype, for instance. However, here I am concerned only with classification into disease types, and stage/grade as well as genetics help mainly to classify cancer according to progression, prognosis, and subtype rather than disease type.

infectious diseases are *not* caused by environmental transmission of the germ. Candidiasis, for instance, is due to overgrowth of the Candida yeast, which is part of our normal microbiome but can proliferate and manifest as an infection due to antibiotic use or immune compromise.

It is a synthetic fact that most infectious diseases are caused by transmission of a specific microbe from without, a general pattern discovered by the germ theorists. In Chapter 2, I called it the Dissemination Hypothesis: disease-causing germs do not arise spontaneously in the host but are introduced into the host from without. Combined with the fact that an infection with a specific germ results from exposure to that same specific germ (there is no transmutation of one species of germ into another), the Dissemination Hypothesis explains why most infectious diseases have a specific universal etiology. Therefore the universal etiology of an infectious disease is partly explained by how we define the disease (as consisting in an infection with a specific germ) but is also partly a brute fact about the way that infections are acquired, a theoretical rather than analytic truth about their causation.

As the history of the germ theory demonstrates, identifying a universal etiology is partly a matter of discovery. The Dissemination Hypothesis was supported by laboratory experiments that refuted the hypothesis of spontaneous generation of germs. Thus the possibility remains of discovering a universal cause for some of our noncommunicable diseases under the constitutive model—the monocausal model is not the only route to universal causation. For instance, a combination of epidemiologic and laboratory research revealed that smoking is a nearly universal cause of small-cell lung cancer, asbestos is a nearly universal cause of pleural mesothelioma, HPV is a nearly universal cause of cervical cancer, and mutations in Rb are a universal cause of familial retinoblastoma (Kumar et al. 2015). Classification is essential to this process of discovering universal causes. Without the proper taxonomic distinctions between malaria and dengue fever, transmission of the plasmodium parasite via mosquito bites is not a universal cause of malaria. Similarly, without just the right analytic distinctions between small-cell carcinoma, pleural mesothelioma, and other malignancies of the lung and chest cavity, smoking is not a nearly universal cause of small-cell lung cancer and asbestos is not a nearly universal cause of pleural mesothelioma. Both discovery and classification have a role to play in the quest for the universal cause.

When universal etiology is discovered through research rather than decreed through definition, it is more often theoretically and

practically worthwhile. First, unlike when stipulating universal etiology via the monocausal model, discovering universal etiology can provide general explanations that are both surprising and add to our understanding. It was surprising to learn that for many infections, all cases are caused by environmental exposure to a germ, and to learn that nearly all cases of pleural mesothelioma are caused by environmental exposure to asbestos. It added to our knowledge of these diseases because these were not analytic truths. Thus to say that people contract malaria because of exposure to the plasmodium parasite and that others develop pleural mesothelioma because of exposure to asbestos is to provide interesting causal explanations for these diseases.

Further, unlike when stipulating universal etiology through the monocausal model, discovering universal etiology can yield practical benefits for public health and clinical medicine. Recall that without any causal selection, multifactorial etiology lends itself to a multipronged preventive strategy in which public health and medicine tackle multiple causal risk factors. This requires having multiple interventions that can effectively manage these risk factors and mustering the money and will to use them all. It requires twice as much effort to eliminate salt and smoking from one's life than to remove just one of these two risk factors for stroke.

In contrast, preventing pleural mesothelioma in a population and in an individual generally requires only avoiding asbestos because asbestos causes virtually all cases of the disease. Genetic risk factors and other industrial minerals like erionite contribute to only a minority of cases (Testa et al. 2011; International Agency for Research on Cancer 2012); avoiding them will consequently prevent only some cases of mesothelioma. To elaborate using the multifactorial theory, these other factors are component causes in only some sufficient causes of mesothelioma, while asbestos is a component cause in nearly all sufficient causes of mesothelioma, or at least nearly all sufficient causes that are responsible for mesothelioma worldwide. Ceteris paribus, if we had to choose to eliminate one exposure or another, we should choose to eliminate the exposure that causes more cases of the disease, and all the better if we eliminate an exposure that causes nearly all cases of the disease. We still need an effective means of eliminating or neutralizing the universal etiology if discovering it is to have any payoff, but this problem is one for etiology generally and not a special problem for universal etiology.

Of course, banning asbestos from buildings will prevent cancers of the pleural lining regardless of how we classify cancers of the thorax. We don't need to define pleural mesothelioma as a cancer of the pleural lining

to reap the benefits of banning asbestos. Yet specific pathology can help in discovering etiological factors, including universal etiology. When a disease has nonspecific pathology, it is more likely to have nonspecific etiology, and the influence of the cause on the occurrence of disease will be diluted in epidemiological data. In contrast, specific pathologies are often more strongly associated with specific—sometimes universal—causes. We are more likely to notice an association between a cause and a pathologically defined human disease when the association is stronger; strength of association is also a consideration commonly used in epidemiology to bolster the conclusion that an exposure is a cause of a disease (Hill 1965). For instance, the definitive epidemiological study establishing that asbestos causes pleural mesothelioma found that nearly all cases of mesothelioma in a series from South Africa's North Western Cape had a history of exposure to asbestos mines in the Asbestos Mountains, while mesothelioma was virtually absent elsewhere in the country (Wagner et al. 1960).

One might object that it is because smoking is a *nonspecific* cause of so many different diseases, including not only coronary artery disease but also disease in other arteries as well as many cancers, that it is such an important etiologic factor for disease prevention, more important than asbestos. But this objection confuses specific causes with specific effects. Indeed, because smoking has nonspecific effects and causes many kinds of disease, avoiding smoking brings great public health and medical benefits. At the same time, if asbestos was not a specific universal cause of mesothelioma, then avoiding asbestos would do less toward preventing this particular disease. While it is often practically advantageous to discover causes that have nonspecific effects (Neal 2019), it is also often advantageous to discover a specific universal cause of a particular disease.

Of several related concepts of specificity canvassed by James Woodward (2010), he suggests that the least important for the idea of exerting fine-grained causal control is the notion that some effect is always brought about by the same cause or causal state, which is similar to the concept of universal etiology. Jacob Neal (2019) goes further and argues that this sort of specificity can actually be undesirable in medicine. Neal provides the example of antiretroviral therapy for HIV, where medicine has multiple drugs within a medication class, each of which can inhibit a particular HIV protein (and thus none of those drugs is used universally by patients to causally control their HIV). However, Woodward and Neal have different contexts in mind. Woodward is mainly interested in the ability of causes to have precise

effects, rather than whether effects have precisely one cause of a certain kind. Meanwhile, Neal's example is one in which the cause is a medical intervention, where it helps to have alternative interventions (for myriad reasons, including differing availability and patient tolerability). When the cause is instead an etiologic factor, and we are instead interested in eliminating its disease outcome as often as possible, then it is harmful to medicine's interests when alternative causes can substitute for that cause.

I noted that the monocausal model can fruitfully organize research on medical prevention by orienting disease categories around specific etiologic factors. For example, because vaccine immunity is generally specific to antigens on particular pathogens, we can potentially develop and deploy an effective vaccine against a disease when that disease is defined according to a particular pathogen. The monocausal model does so (by classifying infectious diseases based on exposure to a particular pathogen)—but so too does the constitutive model (by defining them according to the particular pathogen that is infecting the body). When a specific pathology is associated with a universal etiology, then the constitutive model can match the performance of the monocausal model in research on prevention by making a disease identical to that pathology.

When it comes to research on therapy, specific pathology often performs better than specific etiology because therapeutics intervene on pathology rather than etiology. An etiologic factor will often be nonspecific in the pathology it causes, as illustrated by the example of smoking, which is a nearly universal cause of small-cell lung cancer but also causes myriad other diseases, including other cancers. A particular cancer pathology is often best treated with a unique therapeutic regimen, sometimes involving a targeted cancer treatment, one that interacts with the cancer in a pathology-specific manner. By using the constitutive model to define cancer types or subtypes according to specific pathology, we have developed, studied, and deployed specific therapeutic regimens. Most dramatically, surgery is targeted to the anatomical location of the cancer. But the effectiveness of many aspects of the therapeutic regimen, and even the goals of therapy (e.g., cure versus palliative care), often depend on pathological features of the cancer, including its location, cell type, histopathology (grade), extent of spread (stage), genetics, and gene expression. For instance, targeted cancer therapies such as monoclonal antibodies bind selectively to cancer proteins the same way that the body's antibodies bind selectively to a pathogen's antigen. And just as vaccine-acquired immunity is often specific to a particular pathogen because

the antigen is unique to that pathogen, targeted cancer therapies are often specific to a particular cancer pathology because the cancer protein is unique to that pathology (Chin-Yee 2024).

In summary, many infectious diseases and some cancers do have a universal etiology, but this arises not solely through definition but at least partly through discovery when diseases are defined according to a fruitful specific pathology. Specific pathologies and their universal etiologies can then yield interesting causal explanations and facilitate effective interventions.

Of course, our success in treating and preventing cancers and infections is mixed, which shows that specific pathology and universal etiology are no guarantee of specific and effective treatment and prevention. Moreover, there are many ways of carving up disease categories into specific pathologies that are more or less specific and are characterized through different sciences or levels of description (e.g., anatomical, molecular). So while specific pathology can be a useful classificatory principle, it is not specific enough in its guidance to be the only word on disease classification. In fact, a specific pathology or etiology may be universally associated with a particular kind of disease, and yet our interventions that target that pathology or etiology may be less effective than interventions that are nonspecific in their target or interventions that target a different, nonuniversal etiology. For example, arguably, interventions that strengthen the immune system such as improved nutrition have been historically more impactful in tuberculosis prevention than the BCG tuberculosis vaccine (McKeown 1976a).

Nonetheless, without locating specific pathology, medicine has sometimes struggled to arm itself against diseases. In Chapter 4, I suggest that this includes many chronic noncommunicable diseases and mental disorders. Specific pathology—that great innovation of nineteenth-century modern medicine—could help in furthering our understanding and management of these diseases of the new modern medicine, even if it is not a magic bullet (and does not always yield magic bullet interventions).

## 3.8. Conclusion

The monocausal model is a myth; specific pathology rather than specific etiology was the key conceptual innovation of modern medicine and holds promise for explaining and intervening in disease. It has been proposed that a monocausal model is used to classify infectious diseases according to a

specific necessary etiology and that noncommunicable diseases instead adhere to a multifactorial model in which diseases have multiple and variable etiology. However, the multifactorial model is best viewed as an epidemiological theory of disease etiology representing causal interaction, causal insufficiency, and causal variability, rather than a model of disease classification. Meanwhile, the monocausal model is *not* the model used to classify infectious diseases, nor is it particularly theoretically and practically fruitful as an ideal. Rather, a constitutive model, in which diseases are defined according to the bodily condition of which they consist, describes the classification of both communicable and noncommunicable diseases. The constitutive model can often establish diseases that have a universal etiology by defining them according to a specific pathology. Doing so often provides interesting causal explanations as well as specific and effective medical interventions.

In Chapter 4, I will argue that chronic noncommunicable diseases and mental disorders may pose a challenge for the new modern medicine because they lack a specific pathology. To understand why, we must probe the nature of chronic diseases and mental disorders.

# 4

# The Nature of Chronic Diseases and Mental Disorders


## Summary

During the twentieth century, the epidemiology of disease shifted, as more and more people were diagnosed with and died from noncommunicable diseases, including chronic conditions. Political and metaphysical questions are frequently raised about the legitimacy of certain categories of chronic disease and mental disorder or about whether they are "real" or "natural." But what are chronic diseases and mental disorders, in the first place? What kind of a thing are they? I argue that chronic diseases and mental disorders are generally bodily dispositions, properties that are disposed to manifest clinical findings.

While chronic diseases and mental disorders are a motley collection of maladies, they generally share two features: chronicity and the fact that they are often silent or asymptomatic even while they are typically defined in terms of certain clinical manifestations (the "silent disorder phenomenon"). While chronicity results partly from medicine's historically contingent inability to cure chronic diseases and mental disorders, the silent disorder phenomenon results partly from their nature: chronic diseases and mental disorders are dispositions. Lacking a scientific theory that can illuminate their collective nature, I propose a dispositional theory of chronic diseases that is consistent with how they are diagnosed and understood in medicine. Chronic diseases must be distinguished from their pathology, or the categorical structures or dysfunctions on which they depend, as well as their etiology, pathogenesis, pathophysiology, and manifestations. Meanwhile, the dispositional theory must be defended against alternative metaphysical pictures of chronic disease, including the view that they are fictions, objects, clusters of symptoms/signs/behaviors, or processes or that they include a specific etiology. After ruling out the alternatives, I extend the dispositional theory to mental disorders and argue that mental disorders are dispositions



*The New Modern Medicine*. Jonathan Fuller, Oxford University Press. © Oxford University Press 2025.
DOI: 10.1093/9780190066178.003.0005

toward the symptoms and behaviors that comprise their diagnostic criteria. Finally, I argue that their lack of specific pathology might explain why certain kinds of chronic disease and mental disorder are only weakly projectible; that is, inferences from one sufferer to another might be unreliable in medical explanation, prediction, and intervention.

## 4.1. The New Epidemics

A recurring theme in the chapters of Part 2 of this book is that while acute infectious diseases were the model diseases of nineteenth-century modern medicine, the new modern medicine is greatly oriented toward the care of patients with noncommunicable diseases, the new epidemics. Chapter 2 focused on a family of diseases that are noncommunicable and often acute: cancers. In this chapter, I examine members of two other families of diseases that are also typically noncommunicable but are typically not acute: chronic diseases and mental disorders. I search for a metaphysical theory that provides a general understanding of the nature of chronic diseases and mental disorders.

In Chapter 1, I wrote about epidemiology's classic "epidemiologic transition," a shift in morbidity and mortality from infectious diseases to noncommunicable diseases, especially in developed countries, that was completed in the twentieth century (Omran 2005). Deaths from noncommunicable diseases rushed in to fill the gully left behind by receding epidemics of infectious diseases. The main noncommunicable killers were heart disease—a chronic disease—and cancer. However, other chronic noncommunicable conditions, including chronic respiratory diseases and diabetes, also joined the lists of leading causes of death worldwide.[1] In all, noncommunicable diseases accounted for 74% of deaths globally in 2019 (Global Burden of Disease [GBD] 2020). Meanwhile, in the U.S., rising rates of "deaths of despair" from suicide and substance use after 2000 began to influence and even depress life expectancy, particular for certain social groups (Case and Deaton 2020). In 2019, the worldwide prevalence of mental disorders was 13% (GBD 2020). Although there is some disagreement about just how prevalent and problematic chronic diseases were in the nineteenth

[1] In 2019, the ten leading causes of death in the U.S. were (1) heart disease, (2) cancer, (3) accidents, (4) chronic lower respiratory diseases, (5) stroke, (6), Alzheimer's disease, (7) diabetes mellitus, (8) kidney disease, (9) influenza and pneumonia, and (10) suicide (Xu et al. 2021). COVID-19 joined the list at number 3 in 2020.

century, there is general agreement that chronic diseases are now more common, are a more common cause of death, and are a major healthcare priority globally (Omran 2005; WHO 2005; Weisz 2014).

The prevalence of chronic diseases rose as the aging of the population created an opportunity for people to accumulate chronic diseases, including multiple chronic diseases at once (Hajat and Stein 2018). This phenomenon is foreshadowed in proverbial fashion by a quote commonly attributed to William Osler: "The way to live the longest is to acquire a chronic disease and take good care of it." Soon doctors were writing about epidemics of chronic and noncommunicable diseases (Weisz 2014). However, unlike previous infectious disease epidemics, which were time-limited punctuation marks on a precarious existence, chronic disease epidemics are stable and unremitting, sometimes bringing with them prolonged misery as well as death (though they can also be silent, as we will see).

The World Health Organization (WHO) defines chronic diseases as "diseases of long duration and generally slow progression" (R. Goodman et al. 2013). They currently include cardiovascular diseases like coronary artery disease, hypertension, and heart failure; respiratory diseases like chronic obstructive pulmonary disease and asthma; metabolic diseases like diabetes mellitus and osteoporosis; and gastrointestinal diseases like gastroesophageal reflux disease. Which diseases made the list of typical chronic diseases changed throughout the 1900s, as old chronic infections like tuberculosis and syphilis became curable, while new chronic infections like HIV disease appeared and chronic noncommunicable diseases like diabetes increased in prevalence. Chronic diseases were variously associated with concepts of disability, aging, infirmity, and degeneration throughout the century (Weisz 2014). But a common thread tying chronic diseases together—and one that can change with medical progress—is their incurability.

Their incurability and the difficulty of preventing chronic diseases stand out against the history of the great therapeutic discoveries of modern medicine, especially antimicrobials (including antibiotics) and vaccines. Rather than being cured, chronic diseases are "managed" in the sense that their manifestations are controlled and their progression is halted or delayed (in between acute episodes of "exacerbation" or "decompensation"). These new tasks of medicine often involve the use of long-term medications. Drugs for chronic diseases, especially for long-term primary or secondary prevention, came to dominate prescribing by the end of the 1900s (Greene 2007). Unfortunately, heterogeneity in prognosis and treatment effectiveness is a

common occurrence with chronic diseases, and it is often difficult to make predictions from previous cases of chronic disease with confidence.

In Chapter 6, I will analyze the concepts of risk and risk reduction with which chronic diseases are now intimately associated. In this chapter, my purpose is to understand the nature of chronic diseases. Although chronic diseases differ notably in their causes, characteristics, and effects, there are equally notable features that they share in common. I will discuss two such features in order to motivate my search for a general understanding of chronic diseases. The first feature is their chronicity, their long duration. The second feature is their silence: chronic diseases are often asymptomatic and without overt manifestations, even while they are defined by their manifestations. What are chronic diseases, such that they are chronic and often silent?

These two features—chronicity and silence—are often shared by mental disorders. Mental disorders are equally important to understanding disease and disorder in the new modern medicine. With their rising global incidence, high global prevalence, and high global burden of disease (GBD Mental Disorders Collaborators 2022), mental disorders are also described as epidemics. In the U.S., diagnosed depression is very common, and some surveys suggest it is becoming more common in recent decades (Weinberger et al. 2018; Witters 2023). Moreover, like chronic diseases, mental disorders often behave relatively unpredictably. Because they share important features with chronic diseases, I will extend my analysis of the nature of chronic diseases to mental disorders.

Unlike cancers, about which the somatic mutation theory provides some general insight (Chapter 2), the daunting heterogeneity among chronic diseases and mental disorders means it is unlikely we will find a scientific theory for their pathogenesis or pathophysiology that illuminates their collective nature. Another approach we might consider in trying to get under their skin is concept analysis,[2] the dominant approach to understanding disease in the philosophy of medicine. We could analyze our concept of disease or mental disorder and extract from it any details about the nature of instances of chronic disease or mental disorder. However, ultimately if we want to understand why these conditions are chronic and often silent, we need to better understand the phenomenon itself: the thing to which our diagnoses refer. Thus I will turn to metaphysics in search of an answer to

[2] George Weisz (2014) studies the history of the chronic disease concept in twentieth-century America, Britain, and France, arguing that the importance of chronic disease as a concept hinged on the functions it served in organizing healthcare.

the question *What kind of a thing are chronic diseases and mental disorders?* In particular, are they objects, similar to pathogens? Properties or states, analogous to body mass? Processes, akin to homeostasis? Clusters of signs, symptoms, and clinical features? Or perhaps there is a sense in which they are not real, or are fictions?

In Chapter 3, epidemiology supplied a metaphysical theory of the multifactorial etiology of diseases. Here I develop and defend my own dispositional theory of chronic diseases and mental disorders: chronic diseases and mental disorders are generally dispositional properties of the body. The dispositional theory also helps us understand why chronic diseases and mental disorders are relatively unpredictable. Because a given kind of disposition can depend on a different kind of bodily pathology in different cases, projections about prognosis or treatment from one case to another will often fail.

In section 4.2, I begin the search for a theory of chronic diseases and mental disorders by examining two of their common features: their chronicity and their silence. In section 4.3, I elaborate the dispositional theory through examples and defend it against a slew of metaphysical alternatives. In section 4.4, I extend the dispositional theory to mental disorders. Then in section 4.5, I use the dispositional theory to understand medicine's difficulty in explaining, predicting, and intervening in disease.

## 4.2. In Search of a General Understanding of Chronic Diseases

Chronic diseases are a motley of maladies that include metabolic, cardiovascular, immunological, infectious, and anatomical problems ranging from the mild and well controlled to the disabling and deadly. Despite their differences, is it possible to come to a general understanding of certain interesting features of chronic diseases? For instance, why are chronic diseases chronic? And why are they often silent, without manifesting symptoms or other harms? I will address each of these questions in turn in the next two subsections.

### 4.2.1 Why Are Chronic Diseases Chronic?

In addition to the fact that chronic diseases are often progressive, worsening over time, much of the clinical and public health significance of chronic

diseases can be attributed to their chronic time course. Might something general about the nature of chronic diseases explain why they are chronic? As we will see, chronicity corresponds to a brute fact about most cases of chronic disease: the persistence of some pathology. Yet the proper explanation for why chronic diseases persist has more to do with the state of medicine than with the general nature of chronic disease pathology.

First, it's worth noting the role of epidemiology in recognizing and studying the chronicity of chronic diseases. The chronic course of chronic noncommunicable diseases is best observed in epidemiologic studies with a long-term outlook, including backward-looking (retrospective) and forward-looking (prospective) studies. The investigators of the Framingham Heart Study, a prospective cohort study of heart disease begun in the mid-1900s and continuing today, labeled their study an "epidemiological approach to heart disease" and noted that it was the first long-term study of its kind (T. Dawber et al. 1951). Studying a chronic condition like heart disease required the long-term epidemiological approach because heart disease develops over decades before manifesting in clinically apparent heart ischemia. In fact, the epidemiological approach facilitates the very recognition that conditions like heart disease have a long duration, slow progression, and delayed manifestations.

A couple of clarifications are in order. Typically, when public health authorities mention "chronic diseases" they are referring to disease *types*. A particular case of sinusitis might have a chronic duration, but as a category of disease sinusitis is not considered a chronic disease. (That is why there is a subtype of sinusitis non-redundantly labeled *chronic* sinusitis.) On the other hand, hypertension and osteoarthritis are chronic diseases because the great majority of instances of those diseases are long-lasting. How long is "long-lasting"? Some definitions include a numeric cutoff (commonly, "longer than a year"), but others do not (R. Goodman et al. 2013). The WHO's definition merely states that chronic diseases have a "long duration." While vagueness in the durational cutoff for chronicity leaves room for metaphysically irrelevant factors to determine which diseases get counted as chronic diseases (such as their prevalence in high- versus low-income countries, or their impact on patients and society), that cannot harm my account of the nature of chronic diseases because, as we will see momentarily, I do not assume there is anything metaphysically deep about chronicity.

Despite lack of consensus on precisely how long a disease must endure to be considered a chronic disease, we can still wonder what exactly is doing

the enduring in those diseases that are considered chronic diseases by widespread consensus. On reflection, it is not necessarily the pathogenesis of the disease, the mechanism generating the disease, that is chronic. For example, the process of viral transmission and infection happens relatively quickly in chronic infections like HIV disease and chronic viral hepatitis. Nor is it necessarily the manifestations of the disease, which can be transient and infrequent, especially if the disease is well managed. Rather, in chronic diseases, it is the disease itself that is long-lasting: the core pathology underlying the disease. Not all pathologies encountered in the clinic are long-lasting; for instance, many infections and injuries are not. Therefore, to understand why chronic diseases are chronic, we must explain why it is that chronic disease pathology persists in particular.

There are three ways that a disease can end: in spontaneous recovery, in the death of the patient, or in medical cure. In contrast to their acute cousins, generally chronic diseases do not spontaneously resolve, are generally not fatal in the short term, and are generally incurable. Why do chronic diseases not spontaneously resolve? Sometimes the reason is poorly understood, as with essential hypertension and type 2 diabetes. In other cases, we can pin responsibility on the body's lacking the capacity for repairing certain lesions. Type 1 diabetes mellitus results from the loss of insulin-producing pancreatic islet cells; heart failure can result from damaged and pathologically remodeled heart muscle (after, for instance, a heart attack); and neurodegenerative diseases result from the death of neurons. The adult body has limited ability to regenerate lost endocrine-producing cells of the pancreas, cardiac muscle cells, and neurons in the central nervous system, in contrast to its much greater ability to regenerate epithelial cells in the skin or gut, for instance (Zakrzewski et al. 2019).

Why are chronic diseases generally not fatal in the short term? In some cases, it is because the manifestations of the disease are typically not lethal (as with acid reflux, the characteristic manifestation of gastroesophageal reflux disease, or GERD). In other cases, it is because the potentially deadly manifestations of the disease are well controlled (as with therapeutic insulin preventing hyperglycemic shock, a life-threatening complication of uncontrolled type 1 diabetes). In these latter cases, the disease is not fatal in the short term because of the existence of effective medical treatment.

Finally, why are chronic diseases generally incurable? Many chronic diseases are incurable because they are associated with irreversible anatomical lesions, as with the destroyed islet cells in type 1 diabetes, the

pathologically remodeled heart in heart failure, and the destroyed lung alveoli in emphysema. Yet there are such things as reversible anatomical lesions. Most bone fractures will heal if we cast them. Irreversible anatomical lesions are irreversible because of a lack of reversing medical treatment. Just as the body is unable to endogenously repair the damaged tissue in type 1 diabetes, heart failure, or emphysema, neither are doctors. Other chronic diseases, such as type 2 diabetes and essential hypertension, are incurable because we do not understand enough about the disease to intervene and cure.

In summary, there is very little in the nature of chronic diseases that explains why they are chronic, except that many diseases result from injury to tissues that have limited ability for self-repair. Rather, many potentially deadly chronic diseases are not fatal due to the existence of life-prolonging medical treatment, and many chronic diseases are irreversible or incurable because we lack knowledge of how to reverse or cure them. It does not seem biologically implausible that we might be able to therapeutically reverse or cure chronic lesions at some point in the future. The promise of "regenerative medicine" (Zakrzewski et al. 2019) is to reverse the destruction of pancreatic islets, heart tissue, lung alveoli, and other bodily structures by replacing lost islet cells, cardiac cells, alveolar cells, and other cell types through stem cell therapy. Thus these lesions of chronic disease are not irreversible and incurable in principle but are merely irreversible and incurable at present. To say that certain chronic diseases are irreversible and incurable at present is to make a statement about our current medicine; it is simply to say that currently physicians cannot cure them. Perhaps physicians will be able to cure them in future, perhaps not. Thus it is not the nature of the disease state but the present state of medicine that largely explains why chronic diseases are chronic.

Consequently which diseases are chronic diseases (and which diseases are not) is at least partly historically contingent, dependent on our current medical science. For instance, HIV disease and type 1 diabetes were once more acute than they are now. Those afflicted with HIV disease would typically succumb to the immune-compromising complications of the disease within a few years (Justiz Vaillant and Gulick 2022), while those suffering from type 1 diabetes would only live for a few years at best following diagnosis (Gale 2002). That changed with the development of antiretroviral therapy regimens for HIV disease and insulin for type 1 diabetes, leading to their "transmutation" from acute to chronic diseases (Feudtner 2003).

Conversely, some diseases that were once generally chronic and incurable are no longer. Perhaps the most important examples are tuberculosis and syphilis (Weisz 2014), notorious scourges of the nineteenth century. During the nineteenth century, tuberculosis was a common cause of death, syphilis was a common cause of disability and disfigurement, and both occupied a major place in public consciousness. Tuberculosis and syphilis became curable during the first half of the twentieth century through the advent of antimicrobials, especially the antibiotics streptomycin and penicillin in the 1940s.

Chronicity is a mirror that reflects the new modern medicine back to itself. The growing burden of chronic diseases in the twentieth century is a reminder about the limits of curative medicine. While scientific theories of wound healing and stem cell biology go some way toward explaining why certain chronic diseases persist, medical technology is still required to complete the explanation given that other anatomical lesions that the body cannot fix *can* be repaired by the doctor's hand. Further, as we currently understand them, chronic diseases such as type 2 diabetes, hypertension, HIV disease, GERD, osteoporosis, and asthma are not problems of tissue repair at all. Taken together, these considerations suggest that it is unlikely we will find a general scientific explanation for chronic diseases that explains their chronicity the way that somatic mutation theory explains the pathogenesis of cancers (Chapter 2). The prospects for a general theory of chronic diseases appear dimmer. But before we abandon the search, let us turn to the second of our two motivating questions.

### 4.2.2 Why Are Chronic Diseases Often Silent?

In the 1960s, doctors and public health figures began to speak of a "silent epidemic" of hypertension. While earlier in the century hypertension was diagnosed only when it was acutely symptomatic, in the second half of the century asymptomatic hypertension joined symptomatic hypertension as a diagnosable disease (Greene 2007). The same pattern was repeated for type 2 diabetes and for dyslipidemia, or abnormal blood cholesterol. By the end of the 1900s, nearly all chronic diseases diagnosed by doctors could be silent much or even most of the time. Coronary artery disease and osteoporosis can be silent for years before they erupt into a heart attack or crumble into a fracture, respectively (if they do so at all). Asthma and GERD can be largely

asymptomatic with the help of medication. Many of these diseases can even be diagnosed on laboratory testing or imaging in the absence of symptoms. Because the problem of asymptomatic chronic conditions recurs throughout this chapter, I will give it a name: "the silent disorder phenomenon." What accounts for the silent disorder phenomenon, and can it reveal anything general about the nature of chronic diseases?

We could immediately latch onto the story just told about the history of hypertension, diabetes, and dyslipidemia to provide an explanation that—like the reason I gave for chronicity—is about medicine more than its maladies. Namely, chronic diseases are often silent because epidemiology and medicine recognize chronic diseases that are merely risk factors or precursors for manifestly harmful conditions but that are not themselves manifestly harmful.[3] Moreover, medicine's increasingly precise diagnostic tests and increasingly expansive screening tools are detecting diseases at earlier stages, before they present with manifestations (Aronowitz 2015). And finally, those chronic diseases that are symptomatic can often be medically controlled. So perhaps medical practices and technologies account for the silence of chronic diseases.

While this explanation is convincing, it leaves something unanalyzed about the nature of chronic diseases. While chronic diseases often do not manifest bodily signs or symptoms at a given moment, typically (and paradoxically) they are characterized by their manifestations. Hypertension is characterized by high blood pressure, even when controlled through antihypertensives; diabetes through high blood sugar, even when controlled by insulin or oral antidiabetic agents; coronary artery disease through reduced blood flow to heart muscle, even when asymptomatic; and so on. What can this common feature of chronic diseases reveal about their nature? What kind of a thing are they, such that we understand them primarily through symptoms, signs, or harms that often do not obtain? This question leads us away from scientific theory and toward metaphysical theory.

[3] Examples of such risk factors are dyslipidemia, borderline hypertension, non-insulin-dependent type 2 diabetes, osteoporosis, and class 1 obesity. It is controversial whether they are rightly considered diseases rather than merely risk factors (Schwartz 2008). Before the second half of the 1900s, conditions like hypertension, diabetes, and abnormal blood lipids were diagnosed only when they were symptomatic. In the second half of the century, with the discovery that asymptomatic forms conferred a heightened risk of harms and with the discovery of new medications that could lower these risks (new antihypertensives, oral antidiabetic agents, cholesterol-lowering drugs), the diagnostic borders of these chronic diseases expanded, enveloping asymptomatic forms (Greene 2008). Epidemiology made this pathologization of asymptomatic risk factors possible by providing the research needed to measure risk and pharmacologic risk reduction.

As Jeremy Simon (2010) has observed, metaphysical questions about diseases have received scant attention from philosophers, especially compared to questions about the concept of disease. A recent exception is Simon (2008, 2010, 2011) himself. However, Simon directs more attention to the question of whether diseases are real, which he compares to realist/antirealist debates about theoretical entities in the philosophy of science. Here I am mainly interested in the more Aristotelian question of what kind of a thing a disease is. To what category of being—object, property, process, or other—does a disease belong?

In exploring the nature of chronic diseases, it is important to be clear on our target. Scholars have used "disease" or "disease entity" to refer to a particular diagnostic category or taxon (Severinsen 2001; Whitbeck 1977; Hucklenbroich 2014) or to the textbook model or knowledge of a particular condition (King 1982; Simon 2008). Wondering about the nature of diagnostic categories might mean wondering whether they are natural kinds or artificial constructs. Meanwhile, wondering about the nature of textbook models of disease might mean wondering about whether they are mathematic models or mechanism schema. Although these are interesting questions, in this chapter I will investigate the nature of *instances* of disease. Wondering about the nature of instances of disease involves asking what the (presumably) concrete disease entity is like (e.g., the arthritis in an arthritic patient). Simon (2011) usefully distinguishes disease types (i.e., diagnostic categories) from disease tokens (i.e., instances of disease).

In investigating the nature of chronic diseases, a tempting starting point would be to begin with the concept of a "chronic disease" or "disease" and then deduce the nature of instances of chronic disease from there. To use a far-fetched illustration, if according to the medical literature or philosophy of medicine literature, a disease is an immaterial spirit that has taken up residence in the body, then all instances of chronic disease are instances of a sort of immaterial thing. I previously argued that an approach that started from a disease concept and tried to extract information about the nature of chronic diseases was unlikely to succeed for several reasons (Fuller 2018c). In brief, definitions of chronic disease in the medical literature (R. Goodman et al. 2013) and definitions of disease in the philosophy of medicine literature (Boorse 2011) typically do not reveal what kind of a thing a disease or chronic disease is. Concept analysts have focused mainly on what distinguishes disease from health, normality from abnormality, rather than the metaphysical nature of diseases. That is presumably because the main purpose of a concept of disease is to make exactly this distinction.

An alternative would be to extract an answer to our question from a relevant medical theory. Throughout history, medical theories or systems have offered metaphysical perspectives on disease. Hippocratic humoral theory, which exerted great influence for millennia, was itself a quasi-metaphysical doctrine, the result of applying natural philosophy in the tradition of the Greek Presocratics to the domain of health and disease (Hankinson 2016). Just as the cosmos had four elements, the body, as a microcosmos, had four humors. Humoral theory viewed disease as consisting in dyscrasia, a *state* of the whole body in which the humors were in disharmony, which could not always be attributed to any local defect of the body's parts. While scientific medicine has left behind grand theories of disease, it does currently have theories of families of disease as well as a metaphysical theory of multifactorial etiology, discussed in Chapters 2 and 3, respectively.[4]

However, as I alluded, it seems unlikely that any scientific theory could explain the pathogenesis/pathophysiology of a family of diseases as diverse as chronic diseases, especially considering that one of the only things chronic diseases share in common is that medicine is presently unable to cure them. Moreover, I am simply not aware of any explicit metaphysical theory about the nature of disease or chronic disease in contemporary scientific medicine.

In investigating the nature of chronic diseases, I previously utilized what I called a "bottom-up" or inductive strategy: I analyzed particular types of chronic disease and made inductive generalizations about chronic diseases generally (Fuller 2018c). I sampled common diseases that make up the bulk of what we are typically referring to when we talk about chronic disease—diseases that are unanimously considered chronic diseases and that account for most diagnosed cases of chronic disease in the world. Thus the theory of chronic diseases I developed is broadly representative of instances of chronic disease. The inductive strategy will not necessarily yield a unanimous answer to the question *What are chronic diseases?* (As it turns out, it does not.) However, it can still yield an answer with sufficient scope to serve as an interesting generalization. (As it turns out, it does.)

I chose from a list of twenty chronic conditions compiled by Richard Goodman and colleagues (2013) for the U.S. Department of Health and Human Services (see Table 4.1 for examples). I relied on descriptions of these diseases in the medical literature, making inferences about their

[4] Consistent with the multifactorial theory, the dispositional theory I will propose posits that clinical manifestations result from causal interaction (between a disposition and its stimulus).

**Table 4.1** Examples of Dispositional Chronic Diseases

| Chronic Disease | Description | Characteristic Manifestation | Causal Base |
|---|---|---|---|
| Type 1 Diabetes Mellitus | Glucose intolerance: the inability to physiologically regulate blood glucose, resulting in episodes of sustained hyperglycemia | Hyperglycemia | Insufficient insulin-secreting islet cells of the pancreas |
| Type 2 Diabetes Mellitus | Glucose intolerance | Hyperglycemia | Dysfunction in insulin-responsive cells (and sometimes also in insulin-secreting cells) |
| Chronic Obstructive Pulmonary Disease | A state characterized by airflow limitation that is not fully reversible | Irreversible airflow limitation | Variable (e.g., destruction and enlargement of alveoli in emphysema) |
| Gastroesophageal Reflux Disease | Tendency toward reflux of gastric contents into the esophagus, causing symptoms and/or injury | Reflux of gastric contents into the esophagus | Laxity at esophagogastric junction |
| Hypertension | A state giving rise to high blood pressure | High blood pressure | Variable (e.g., high levels of aldosterone in hypertension secondary to Conn syndrome) |
| Coronary Artery Disease or Ischemic Heart Disease | Tendency toward episodes of myocardial ischemia: insufficient blood supply to heart muscle | Myocardial ischemia | Narrowness in coronary arteries due to atherosclerotic plaque |
| Asthma | Airway hyperresponsiveness: tendency toward a physiological state giving rise to airflow obstruction and a constellation of symptoms (wheezing, shortness of breath, coughing) | Reversible airflow obstruction, or wheezing, dyspnea, coughing | Poorly understood, but may include chronic inflammation of the respiratory mucosa |
| (Left Ventricular) Heart Failure | Left ventricular dysfunction, resulting in circulatory congestion and signs and symptoms of circulatory congestion (e.g., shortness of breath, edema) | Low cardiac output, or signs and symptoms of circulatory congestion (e.g., shortness of breath, edema) | Variable (e.g., left ventricular remodeling in more advanced, progressive heart failure) |
| Osteoporosis | Decreased bone strength due to low bone density | Bone fractures | Low bone density |

*Source:* Kasper et al. (2015).

nature from the way they are defined, discussed, and diagnosed.[5] Goodman and colleagues listed several mental disorders, but I will consider mental disorders separately in section 4.4. While Goodman and colleagues also list cancer as one of their twenty chronic conditions, I will not consider cancer for several reasons. First, I adopted a different approach to understanding cancer and already carried it through in Chapter 2. Second, according to that analysis, cancer is a family of diseases rather than a species of disease at the same level of taxonomic organization as, say, asthma. There are certainly types of cancer that are recognized as having a chronic time course, such chronic lymphocytic leukemia, as well as cases of cancer that are slowly progressive, such as many cases of prostate cancer. Nonetheless, here I will focus on conditions that are more paradigmatic examples of chronic diseases.

Expanding a list of candidates from Simon (2011), I considered whether each of the chronic diseases I surveyed was an object, a bundle of signs and symptoms, a property or state, a process, or a fiction. I further examined whether a chronic disease was one of these aforementioned entities *plus* a specific cause, given the presumed importance of specific etiology in distinguishing diseases (discussed in Chapter 3). While objects, properties, and processes are the three categories most commonly considered potentially fundamental in metaphysics, I will not weigh in on which of these nine categories are truly fundamental. Some of these categories—such as "bundles of signs and symptoms"—are not basic categories in metaphysics at all. I chose these nine candidates because all of them have been vigorously defended at some point in history by physicians or medical commentators.[6]

As a preview of the next section, a survey of common chronic diseases reveals that the vast majority of them are bodily properties. Moreover, most chronic diseases, especially noncommunicable chronic diseases, are more

[5] There are various sources that I could consult and that would provide a very different perspective on chronic diseases, including patient narratives. While patient narratives are valuable, they would mainly capture a phenomenon that is not my current target: namely, chronic *illness*, the first-person experience of chronic disease (Carel 2016). Here I am interested in the referents of chronic disease diagnoses (type 1 diabetes, type 2 diabetes, etc.)—chronic diseases as scientific theoretical entities.

[6] Historians often distinguish "ontological" conceptions of disease from "physiological" conceptions (Rather 1959; Canguilhem 1966; Duffin 2005). Often the ontological conception is taken to be the view that diseases are distinct from the patient (not a proper part of the patient), while the physiological conception holds that diseases are continuous with the patient (a proper part of the patient). On the dispositional theory I propose, chronic diseases are properties of the patient's body and are thus part of the patient, as the physiological conception would have it. However, "the ontological conception" is also sometimes used to refer to the idea that diseases come grouped in species or even that species of disease have an independent existence, in contrast to the idea that there are no disease species, only diseased patients. As noted earlier, my focus is the nature of instances of disease rather than disease categories.

specifically dispositions.[7] Metaphysicians usually recognize two different kinds of properties: dispositional[8] (or realizable) properties and manifest (or categorical) properties. Dispositional properties are said to realize/manifest certain processes or effects on certain occasions. For instance, *solubility* is the disposition of a solute to dissolve, while *fragility* is the disposition of an object to break. Usually (though not always; Vetter 2014) dispositions are also analyzed according to their stimulus or triggering conditions (Lewis 1997), which cause the disposition to manifest. For instance, solubility is the disposition of a solute to dissolve *when the solute encounters an appropriate solvent*, while fragility is the disposition of an object to break *due to stress or impact*. It is usually (though not always; Friend 2021) thought that dispositional properties can be present when their characteristic manifestation is absent, in which case it is said that the dispositional property is present but unmanifested (Molnar 2006). A solute is soluble (in some solvent) whether or not it is dissolved, and a porcelain plate is fragile regardless of whether it has broken. Everyday human dispositions include the conductivity of the heart (a disposition toward electrochemical impulse transmission) and personality traits (dispositions toward certain emotions or behaviors).

In comparison, manifest properties include our more passive and prototypical examples of properties, like shape, size, and density, which are also possessed by people and their body parts. Manifest properties are always fully realized whenever they are present. Whenever a person is tall, their substantial height is manifest.

The dispositional nature of chronic diseases helps explain why they are so often silent even while they are characterized by their manifestations: chronic disease dispositions can be present without manifesting their characteristic manifestation. In the following sections, this will be my general interpretation of the silent disorder phenomenon. There are two general reasons why a chronic disease disposition might not be manifesting. First, it could be that its stimulus conditions haven't triggered the condition (e.g., the asthmatic has not encountered the triggers that would lead to an asthma attack, the

[7] A volume edited by Anjum and colleagues (2020b) contains various proposals for how the theory of causal dispositionalism can improve our understanding of causality in the clinic.

[8] Authors sometimes contrast dispositions with other kinds of realizable properties such as tendencies and propensities (Jansen 2007), but here I will use "disposition" as a broad term that encompasses all kinds of realizable properties. In Chapter 6, I will discuss an interpretation of individual risks in medicine that sees them as propensities, which are chancy realizable properties.

individual with osteoporosis has not met with the fall that would cause a fracture). Second, it could be that its expression is being masked by a treatment. In dispositions talk, a "mask" or an "antidote" blocks the manifestation of a disposition in the presence of its stimulus condition, as when protective padding blocks the fragile object from breaking after it has been dropped (Bird 1998). Many medical treatments work by masking a disposition: inhalers blocks asthma from manifesting in an asthma attack even in the presence of triggers; insulin blocks type 1 diabetes from manifesting in hyperglycemia even after a meal. By avoiding stimulus conditions and seeking antidotes, one could have a chronic disease that is asymptomatic for long periods of time, even when it is defined according to its symptoms or manifestations.

While the chronicity of chronic diseases reflects mainly the historically contingent limits of curative medicine, the fact that chronic diseases are commonly silent even though they are characterized by certain overt manifestations reveals something about the metaphysics of chronic diseases. Namely, chronic diseases are generally dispositions. In the next section, I will further develop the dispositional theory of chronic diseases and defend it against eight alternatives.

## 4.3. Chronic Diseases Are Dispositions

In section 4.3.1, I will elaborate on and illustrate the dispositional theory of chronic diseases through common examples. In section 4.3.2, I will defend the dispositional theory against several rival theories.

### 4.3.1 A Dispositional Theory of Chronic Diseases

With some exceptions,[9] chronic diseases are bodily properties. Take these definitions of osteoporosis and COPD. "Osteoporosis is defined as a reduction in the *strength* of bone that leads to an increased risk of fractures" (Lindsay and Cosman 2015, 2488, my emphasis). Strength is a property of bone. Similarly: "Chronic obstructive pulmonary disease (COPD) is defined as a disease *state* characterized by airflow limitation that is not fully

[9] Perhaps the most important exception is atherosclerosis, a *process* in which atherosclerotic plaque accumulates in the walls of arteries. Atherosclerosis is often considered a disease (R. Ross 1999; Hansson 2005; Herrington et al. 2016).

reversible" (Reilly et al. 2015, 1700, my emphasis). COPD is not an object limiting airflow, nor a pathogenic process resulting in airflow limitation, nor any symptoms or signs of airflow limitation, but a "state" giving rise to limited airflow. Likewise, diabetes mellitus is an intolerance to glucose, asthma is airway hyperresponsiveness, and osteoarthritis is pathological joint structure.

What kind of properties are these chronic diseases? Let's start with the last disease on the list immediately above: osteoarthritis. Osteoarthritis is a manifest or categorical property. It is present just when its characteristic manifestation is present. By "characteristic manifestation" I mean the property or process that defines the disease. Osteoarthritis is defined by structural abnormalities, especially cartilage deficiency. Whenever a joint has osteoarthritis, it has some manifest abnormalities of joint structure that are present so long as the osteoarthritis is present. Similarly, chronic infectious diseases such as HIV disease and chronic viral hepatitis are defined as infections with a virus, and so long as someone has one of these infectious diseases, some of their body's cells are infected with the characteristic virus. This suggests that these diseases are identical with their characteristic manifest property.

However, among noncommunicable chronic diseases, osteoarthritis is an exception. Most of the common noncommunicable chronic diseases—osteoporosis, COPD, diabetes, asthma, GERD, coronary artery disease, heart failure—are dispositional properties (Table 4.1). When a disease is a dispositional property, its characteristic manifestation is the process or property toward which the disease is disposed. Thus it is not the case that the characteristic manifestation is always present when the disease is present. For example, diabetes is often present when its characteristic manifestation (hyperglycemia) is absent due to effective therapeutic control of blood sugar levels (the silent disorder phenomenon). Diabetes is a *disposition toward* hyperglycemia rather than the hyperglycemic state itself. Similarly, osteoporosis is a disposition toward bone fractures, COPD is a disposition toward irreversible airflow limitation, asthma is a disposition toward reversible airflow limitation, GERD is a disposition toward reflux of stomach contents into the esophagus, and coronary artery disease is a disposition toward insufficient blood flow to cardiac muscle.[10]

[10] I won't systematically investigate whether the dispositional theory applies well to acute diseases. However, if the examples of HIV disease and chronic viral hepatitis give us any clue, then it may be that infectious diseases—both acute and chronic—are generally manifest properties (specifically, the state of being infected with a pathogen).

Turning to heart failure, another common chronic disease, Nicholas Binney (2018) argues that the nature of the heart dysfunction in heart failure is incompletely understood and contested. From this, Binney argues that the function of the heart is not obvious, which raises problems with dysfunction-requiring accounts of the concept of disease that use heart failure as an obvious case. Setting aside the matter of whether Binney's conclusion is warranted by his arguments, the dispositional theory allows us to bracket difficult questions about the function of the heart and its dysfunction in heart failure (and about our concept of disease) by offering an account of the nature of heart failure that could help cardiologists usefully refine the way that the disease is defined.

There are a couple of ways that heart failure is typically defined. To use Niels Gadsbøll and colleagues' (1989) distinction, heart failure is considered either a "disease state" or a "clinical syndrome."[11] For example, the 2012 European Society of Cardiology guidelines for the diagnosis and treatment of heart failure define it as a dispositional state: "Heart failure can be defined as an abnormality of cardiac structure or function leading to failure of the heart to deliver oxygen at a rate commensurate with the requirements of the metabolizing tissues, despite normal filling pressures (or only at the expense of increased filling pressures)" (McMurray et al. 2012, 1792). Meanwhile, the updated 2016 guidelines define heart failure as a syndrome: "HF [heart failure] is a clinical syndrome characterized by typical symptoms (e.g. breathlessness, ankle swelling and fatigue) that may be accompanied by signs (e.g. elevated jugular venous pressure, pulmonary crackles and peripheral oedema) caused by a structural and/or functional cardiac abnormality, resulting in a reduced cardiac output and/or elevated intracardiac pressures at rest or during stress" (Ponikowski et al. 2016, 2136).

A virtue of a symptom-based definition of heart failure is that it allows doctors to set aside important uncertainties over the pathophysiology of heart failure by defining the disease "clinically," according to readily observable clinical findings. However, characterizing heart failure as a syndrome fails to respect a widespread diagnostic practice. Many patients are able to ameliorate the signs and symptoms of heart failure using a combination of diet and medication. The syndrome disappears, at least for a time, and yet

[11] Asthma is similarly defined either as a hyperresponsive *state* of the airways giving rise to variable airflow limitation (Bateman et al. 2008) or as a clinical *syndrome* (wheezing, dyspnea, coughing) resulting from this physiological state (Barnes 2012).

the patient retains the diagnosis of heart failure, which strongly suggests that heart failure—at least to the everyday clinician—is not the syndrome itself. The dispositional theory offers a means of tweaking the European Society's symptom-based diagnosis, retaining the virtue of being agnostic about pathophysiology while simultaneously respecting everyday diagnostic practice. Specifically, the dispositional theory can construe heart failure as a disposition toward the characteristic manifest syndrome, such that heart failure remains even when its symptoms are masked through lifestyle or medical intervention.

A different dispositional account of heart failure is also feasible, one in which heart failure is a disposition toward physiological manifestations, more consistent with the European Society's 2012 definition. For instance, cardiologists could define heart failure as a disposition toward reduced cardiac output. This dispositional definition can deftly handle a problem that confounds other attempts to physiologically characterize heart failure: the fact that in many cases of heart failure, cardiac output is preserved (Binney 2018). Again, a disposition can be present but not manifest. The reason that heart failure often does not manifest in reduced cardiac output is because the disposition is masked by increased filling pressure, which compensates for reduced pumping efficiency (Kemp and Conte 2012).[12] This more physiological definition of heart failure can remain agnostic as to the relative importance of mechanical (Kemp and Conte 2012) versus neurohumoral (Haydock and Cowie 2010) mechanisms of cardiac congestion as well as the incompletely understood mechanisms of progression of heart failure, which makes it a viable alternative to the symptom-based dispositional account. In either case, the fact that symptoms of congestion or reduced cardiac output are not always manifest (the silent disorder phenomenon) is demystified when we account for the dispositional nature of heart failure.

Two further features of dispositions are worth discussing. First, it is often held that a disposition is an intrinsic property of its bearer (Lewis 1997). Even if not true of all dispositions (McKitrick 2003), it seems to be true of chronic diseases, which travel with the diseased person through different external contexts. Many chronic diseases are specifically intrinsic properties of parts of the body: osteoporosis is a property of bone, COPD and asthma are

[12] However, the compensatory mechanism(s) that maintain cardiac output in heart failure *do* manifest in signs and symptoms of circulatory congestion, which is why heart failure often calls attention to itself at a time when cardiac output is preserved.

properties of the airways, GERD is a property of the upper digestive tract, coronary artery disease is a property of the coronary arteries, and heart failure is a property of the heart. However, other chronic diseases may not be localizable to an anatomical part, for example, autoimmune or metabolic diseases such as type 2 diabetes. What determines whether a chronic disease is localizable to a part of the body is the localizability of its "causal base."

Second, many metaphysicians have argued that at least some dispositional properties have a causal base and that this base is a categorical property. The disposition depends for its existence on its causal base. Again, even if not true of all dispositions,[13] chronic disease dispositions do seemingly depend on some concurrent bodily pathology. For example, osteoporosis depends on low bone density, COPD depends on anatomical changes in the lungs, GERD depends on an abnormal esophagogastric junction, coronary artery disease depends on narrowed coronary arteries due to atherosclerotic plaque, and heart failure depends on structural or muscular heart abnormalities. These causal bases are themselves localized to body parts. In contrast, type 2 diabetes depends on abnormalities in insulin-responsive cells located throughout the body (and sometimes also on abnormalities in insulin-secreting cells).

The nature of the dependency of dispositions on their causal base is controversial. Some authors argue that at least those dispositions that have a categorical causal base are identical with that categorical base (D. Armstrong et al. 1996); it is merely that we pick out this categorical property using a disposition ascription. On this view, heart failure is whatever manifest properties of the heart cause symptoms of circulatory congestion or low cardiac output given the right stimulus conditions (and in the absence of masking conditions). If this identity thesis holds true of chronic diseases generally, then the dispositional status of chronic diseases is about how medicine characterizes them rather than their independent nature. Consequently, when I argued in the previous section that their dispositional nature explains the silent disorder phenomenon, this would partly be a statement about medical practices after all: we *choose* to identify the disease using disposition talk. We could just as soon identify chronic diseases through their categorical base.

[13] For instance, some dispositions (perhaps certain fundamental physical properties) may be "bare dispositions" that lack a categorical base (Mumford 2006).

For instance, the categorical base for osteoporosis (defined as low bone strength, a disposition toward fracturing) is always low bone density (a categorical property). And sometimes medicine defines osteoporosis according to this categorical base rather than the disposition: "The World Health Organization (WHO) operationally defines osteoporosis as a bone density that falls 2.5 standard deviations (SD) below the mean for young healthy adults of the same sex" (Lindsay and Cosman 2015, 2488). If low bone density as a categorical property is identical to low bone strength as a dispositional property,[14] then the difference between defining osteoporosis as low bone density versus low bone strength rests in how we refer to the selfsame property, namely, through its structural basis or its characteristic effects.[15]

However, as I will discuss in section 4.5, some dispositionally defined types of chronic disease depend on different categorical bases in different token instances; for example, heart failure can sometimes depend on thickening of the wall of the left ventricle, sometimes on narrowing of the left ventricular outlet, and sometimes on other pathologies. So we cannot alternate freely between a dispositional definition of heart failure and an equivalent categorical definition; the two definitions are not coextensive. This is a brute fact about the pathological basis of heart failure once we define the disease dispositionally. We could of course split heart failure into pathologically defined categories. But for other chronic diseases about which our understanding of the pathology is limited, such as type 2 diabetes, this option is not available to us. So even a token identity theorist about chronic disease dispositions and their causal bases would have to concede that dispositional characterization is not always simply a choice among equally viable dispositional and categorical ways of characterizing a type of disease.

Sander Werkhoven (2019) proposes a dispositional theory of health. For Werkhoven, "Health is the ratio of a living organism's dispositional set compared to the maximum dispositional set belonging to its reference class" (934). The "dispositional set" is simply the set of dispositions possessed by an organism. Correspondingly, disease involves the *loss of dispositions* for the

[14] One good reason to accept token identity of dispositions with their causal bases in examples like osteoporosis is that we can assign the disposition's role (its manifesting in bone fractures given certain stimulus conditions) to the categorical base (low bone density). Once we do so, there's no work left for some nonidentical disposition token to do, and no strong reason to commit to its existence.

[15] Implicit in the dispositional definition of osteoporosis is that it depends on low bone density as its categorical base, which distinguishes osteoporosis from other diseases of decreased bone strength that depend on different bone pathology, such as osteogenesis imperfecta (Forlino et al. 2011).

organism, reducing the organism's dispositional set and thus its health. At first glance, this view of disease is consistent with my own because in gaining a disease disposition, an organism could also lose a healthy disposition—there is no contradiction in maintaining that acquiring a chronic disease involves both the loss and gain of dispositions. In fact, chronic diseases are often the mirror image of the healthy dispositions that they replace. For example, heart failure replaces a disposition toward adequate cardiac output with a disposition toward inadequate cardiac output, and osteoporosis replaces a disposition to resist fracturing with a disposition toward fracturing.[16]

Before considering competitor theories of chronic disease in the next subsection, we need in place a few distinctions. We should distinguish among a chronic disease, its pathogenesis, its etiology, its pathology, its pathophysiology, and its manifestations. In determining the nature of the disease, we must be careful not to confuse the disease with these other entities. Let's start with pathogenesis. As the name suggests, pathogenesis describes the genesis of a disease. Turning to some of the examples of diseases with which we have been working, the pathogenesis of type 1 diabetes is the autoimmune *destruction* of insulin-producing islet cells of the pancreas, resulting in a lack of insulin secretion. On the other hand, the pathogenesis of osteoporosis is disordered bone *remodeling*. And the pathogenesis of coronary artery disease, named atherosclerosis, is the *building up* of atherosclerotic plaque in the walls of coronary arteries. Destruction, remodeling, and "building up" all describe *processes* or change. These examples are broadly illustrative: a disease's pathogenesis is whatever process generates the disease. Elsewhere (Fuller 2022b), I argue that the processes that make up a disease's pathogenesis are more specifically etiological mechanisms (see also Dammann (2020)). In contrast, the etiology of a disease typically refers to an upstream cause, often a cause that isn't intrinsic to the body, such as an environmental exposure or an aspect of the individual's lifestyle or behavior.

Moving on, the pathology of a disease is the categorial bodily basis for the presence of disease, what I have been calling the "causal base" of a chronic

[16] That said, it is odd to think that acquiring a chronic disease, as a patently unhealthy disposition, contributes to health (even if the gain of that disposition is offset by the loss of other healthy dispositions). So ultimately my account of chronic disease suggests that a dispositional theory of health like Werkhoven's should distinguish among health-increasing dispositions on the one hand, and health-decreasing dispositions on the other hand. Since Boorse (1977), concepts of health and disease proposed by philosophers have commonly invoked function and dysfunction. While popular accounts of function rely on dispositions, the role of dispositions in dysfunction is controversial even for these accounts (Conley 2023).

disease disposition. It is produced by the disease's pathogenesis. Low islet cell mass, the pathology of type 1 diabetes, is produced by the autoimmune destruction of islet cells; low bone density, the pathology of osteoporosis, is produced by disordered bone remodeling; and atherosclerotic plaque that causes narrowing of a coronary artery, the pathology of coronary artery disease, is produced by atherosclerosis. While in medicine the term "pathology" is sometimes reserved for that which is grossly observable (e.g., on autopsy), I will include under "pathology" lesions that are not observable with the naked eye (e.g., cellular pathology that is observable only under the microscope, as well as molecular pathology that is detectable only through molecular assays).

Next, the pathophysiology of a disease consists in the mechanisms linking the disease or its pathology to its manifestations. For example, the pathophysiology of type 1 diabetes consists of mechanisms resulting from lack of insulin, such as glycolysis (the breakdown of fat), as well as those resulting from high blood glucose, such as diuresis (excessive urine production). And the pathophysiology of coronary artery disease consists of mechanisms of cardiac muscle ischemia (insufficient blood flow), including the reduced blood flow in a narrowed artery associated with angina as well as obstructed blood flow in an occluded artery associated with myocardial infarction. Finally, the manifestations of disease are its effects. We already encountered one kind of manifestation that I named the "characteristic manifestation" of a disease. The characteristic manifestation of a dispositional disease is the defining effect toward which the disease is disposed. Similarly, other manifestations, including observable signs and symptoms of disease, result from the disease's pathophysiology.[17][18]

[17] Caroline Whitbeck (1977) argues against this view that the manifestations of a disease are correctly conceived as effects of the disease. In her view, "the signs and symptoms which are identifying marks of a disease are not properly understood as effects of the disease process but states of affairs that occur within it" (634). Whitbeck argues that signs and symptoms are parts of the disease. A whole cannot cause its parts, and thus a disease cannot cause its signs and symptoms. However, Whitbeck's argument relies on a different disease metaphysics than the one I have proposed. As illustrated by the above quote, Whitbeck views diseases as processes—processes of which signs and symptoms are a part. I will argue against this picture in section 4.3.2, at least as an image of *chronic* diseases. However, Kristina Engelhard (2021) argues for a partial identity between dispositions and their manifestations.

[18] At first glance, hypertension might look like a counterexample to my generalization that signs and symptoms are effects of the disease. Hypertension is often considered a disease (Greene 2008; Schwartz 2008; Aronowitz 2015) and is also commonly identified with high blood pressure, a sign. The dispositional theory helps disambiguate hypertension by distinguishing between hypertension as high blood pressure and hypertension as the chronic disposition manifesting in high blood pressure. Dyslipidemia is another chronic condition that suffers this ambiguity: it can refer either to an abnormal lipid profile or to the disposition manifesting in an abnormal lipid profile.

In summary, chronic diseases are generally dispositions, their causal base consists in categorical pathology of the body or its parts, their pathogenesis is the process or etiological mechanism generating the disease, their manifestations are effects of the disease, and their pathophysiology are the mechanisms linking the disease with its manifestations. As we have already seen, the dispositional theory is not the only theory of diseases on offer. In the next section, I defend my view against a slew of competitors.

### 4.3.2 Rival Theories of Disease

Earlier, I mentioned several alternative answers to the question *What kind of a thing are chronic diseases?*: fictions; objects; bundles of signs and symptoms; processes; and any of the previous three entities—or a property—*plus* a specific cause. The difficulty of answering the question is due in part to the fact that pathological entities of these kinds usually coexist with the disease; there are signs and symptoms, diseased body parts, mechanisms of disease, and etiological factors. The challenge is to determine which entity at the pathological potluck is the actual disease. I will now consider each of the alternatives in turn.

First, there is the fictionalist possibility. One can imagine several versions of this antirealist position (none of which denies that patients diagnosed with chronic disease have real symptoms and suffering). Perhaps our disease diagnoses do not really refer. Perhaps when physicians diagnose a patient as "having arthritis" or "having hypertension," and when they talk about treating arthritis or hypertension, they are being metaphorical rather than literal. Naming the disease could merely be a way of classifying patients who are similar in certain respects. Diagnosing a patient "with" a particular disease might really mean classifying a person as a particular type of patient; treating the disease might be medical-speak for treating that person as a member of that type. However, once we have granted that there are certain respects in which a particular class of patients are similar (in virtue of which we class them together), it is more reasonable to believe that disease talk that appears to be about some of these respects *really is* about these respects of similarity. If osteoarthritic patients really do have certain joint similarities, it is strange to think that when a doctor tells an osteoarthritic patient that they "have arthritis in their joint," what they really mean is that the patient is of the osteoarthritis type.

Another version of fictionalism might maintain that diseases are ideas: when textbooks and physicians describe a disease, they are describing something abstract, not something concrete. This interpretation might explain why two patients with the same disease diagnosis can have unique pathologies and clinical presentations, while a textbook description of the disease depicts a single entity: the disease entity being described is an abstraction from the unique cases. On the surface, Simon appears to propose such a view when he claims, "Ontologically, a disease is an abstract entity" (2008, 362). Yet when an abstraction abstracts away the features that make individual cases unique, it leaves shared features represented. Once we accept that abstract textbook entities represent shared concrete pathology, it is a mistake to think that the abstract entity is the disease token. The disease is what is being represented (the patient pathology), not the representation. As evidence for this claim, diseases cause clinical manifestations: signs, symptoms, and other findings. Concrete pathology can certainly cause clinical manifestations, but abstract representations cannot (at least not directly). The abstract representation is best conceived as a *model of* the disease. On closer reading, Simon seems to support this interpretation. Simon argues that medical textbooks include a model—the "abstract entity" he mentions—and that this model "will represent a (constructively) real disease" (363).[19]

Returning to our list of rival theories, among the seven realist alternatives are the four "cause+" possibilities. These alternatives hold that a disease is an object or syndrome or process or property *plus* a specific etiologic factor. We can dispense with these four possibilities for two reasons. First, as I discussed in the previous chapter, most noncommunicable diseases do not have a specific universal etiology. But second, even diseases that do have a specific etiology do not consist partly in their specific etiological factor. Otherwise we would not describe that factor as the "etiology"—meaning cause—of the disease because doing so would imply that diseases contain their own cause. Smoking is a cause found in nearly every case of small-cell lung cancer. Yet doctors do not say that smoking is one component of the disease that causes other components; they say that smoking causes lung cancer (Alberg et al. 2013). Doctors also tend to view diseases as occurring within the patient and not extending beyond the skin to include the individual's behaviors and

[19] Simon (2008) argues that the disease is "constructively real" because his account adapts Ronald Giere's (1999) constructive realism.

environment, yet etiological factors commonly consist of behaviors and environmental exposures. Thus diseases do not include their etiology.

Another alternative to consider only briefly is whether chronic diseases are objects. This conception of diseases was more prominent in earlier times, when diseases were sometimes thought to be afflictions that invaded the body—spirits or miasmas. Diseases are now conceived as *conditions of* the body. When a person has a disease, there is something wrong with the body itself. With the advance of osteoporosis, bones transition to diseased bones, and with heart failure, a heart transitions to a diseased heart. The disease that afflicts the person with osteoporosis or heart failure is not an object—it is not the diseased bone or the diseased heart itself.

The view that diseases are bundles of signs and symptoms was also historically popular, beginning especially with the physician Thomas Sydenham, who in the seventeenth century distinguished types of disease based on their observable natural history. It has also been defended in more recent times. Lester King writes, "When most people talk of disease, they have in mind some particular condition like diabetes or pneumonia, or peptic ulcer. In this more limited sense, the term refers to a *pattern* of factors which somehow hang together and recur, more or less the same, in successive individuals. Thus, pain in the right lower quadrant of the abdomen, with nausea, vomiting, a fever, and a high white count, spell out the features of acute appendicitis" (1954, 197, my emphasis). L. J. Rather claims that the "clinical entity . . . is a discrete, orderly *pattern* of events"; these events are "phenomena observed by clinicians" (1959, 366, my emphasis). King and Rather seem to be talking about disease *types*. After all, patterns are ideal entities that we abstract from similar cases. The disease tokens that these patterns represent are instead "clusters" (King 1954, 198) of observable clinical features.

Rather uses an example to argue for this picture of disease. Sydenham's conception of gout—like his conception of all diseases—was based on the clinical picture: gout was essentially pain in the great toe or ankle that followed a characteristic clinical course. Rather argues that the clinical entity gout has remained intact since Sydenham's day, despite the advance of physiology and biochemistry. For Rather, the resilience of this disease type despite the changing scientific understanding of the body is evidence that diseases are Sydenham's congeries of clinical features.

Yet Rather's argument rests on a mistaken contemporary conception of gout. Though there is still a diagnostic category labeled "gout" today, it is not the case that our conception of gout has remained unchanged. Sydenham

would classify two presentations of ankle pain that both fit his description as two cases of gout, but today we would classify these two cases differently if they had a different biochemical basis. We would classify one case as gout if there was a buildup of urate crystals in the joint, and the other case as pseudogout if there was a buildup of calcium pyrophosphate crystals (Sidari and Hill 2018). That is because we no longer consider gout to be a particular clinical picture but now consider it to be a property of a joint that has been infiltrated by urate crystals. Medicine underwent a shift in the nineteenth century from a picture of diseases as bundles of clinical features to one of diseases as pathologies (Rosenberg 2002).

An additional argument against the view that chronic diseases in particular are bundles of signs and symptoms leverages the silent disorder phenomenon. Chronic diseases are often asymptomatic, and thus cannot be a bundle of signs and symptoms. Finally, constellations of signs and symptoms are today called "syndromes" rather than "diseases." The two categories of syndrome and disease are generally considered mutually exclusive; thus chronic *diseases* are not clusters of signs and symptoms. The fact that many important chronic maladies are symptoms or syndromes explains why organizations often refer to them collectively as "chronic conditions" rather than "chronic diseases" (WHO 2005; Nasmith et al. 2010; R. Goodman et al. 2013).

The final alternative that I will consider is whether diseases are processes, a view that has several supporters.[20] According to Whitbeck, "The distinction is commonly made between diseases and impairments on the basis of the contrast between static conditions and processes. Static abnormal conditions, like hare lip [cleft lip] say, contrast with diseases in that diseases follow a characteristic course. Thus, chronic disease is defined as one 'of slow progress and long continuance' not as a permanent or enduring state" (1977, 625). Whitbeck concludes that diseases are complex processes rather than states. Others follow Whitbeck in distinguishing static "impairments" or other states from dynamic diseases (Reznek 1987; Nordenfelt 1995; Williams 2007). For example, Olaf Dammann (2020) defends a hypothesis about the etiology of autism in which it develops over time (during pregnancy and after birth) through neuro-immunological and neurodevelopmental processes, and thus it is better to think of autism as an ongoing process rather than a static disorder.

[20] For instance, Whitbeck (1977), Reznek (1987), Nordenfelt (1995), and Williams (2007).

However, the fact that diseases in some sense "follow a characteristic course" (to borrow Whitbeck's expression) does not entail that diseases *are* that course. Chronic diseases are produced by a pathogenic process. That pathogenic process may continue to cause changes even as the disease qua disposition continues to stably exist. In type 2 diabetes, the disposition toward hyperglycemia endures even as insulin resistance and insulin insufficiency worsen. What changes is the severity of the disease. For instance, the disease eventually manifests in hyperglycemia even with certain therapies that previously could mask the disease's realization. Further, the disease itself can produce changes in the body, and these changes are often recognized as part of the characteristic course of the disease. For example, hyperglycemia produces vascular changes in the body. A dispositional theory of chronic diseases is fully consistent with the slow progress of chronic disease pathology.

In contrast, a process view has trouble accounting for important characteristics of several chronic diseases. In general, there are two notable processes taking place in patients with chronic disease, either of which might plausibly be the disease. The first is the progression, course, or natural history of the disease pathology, including its pathogenesis. If type 2 diabetes were identical to its course, the disease would then consist in the following sequence: development of insulin resistance, followed by development of insulin insufficiency (both of which worsen over time), with the progression of complications like retinopathy and nephropathy (which also multiply and worsen over time). However, some chronic diseases have virtually no progression past a certain point; progress ends, yet the disease endures. GERD is a good example. In patients with GERD whose heartburn symptoms are well controlled, the disease course is generally non-progressive; symptom severity might fluctuate from day to day but typically does not worsen from year to year. Atrial fibrillation, another chronic disease, similarly does not generally progress once it develops. If the progress of a disease ends but the disease continues, then the disease must not be the progress. At the opposite end, early in its natural history, a disease is often not yet occurring, and thus identifying the disease with its natural history would lead to a premature disease diagnosis. For example, early in the natural history of type 1 diabetes, when the immune system just begins attacking the pancreas, enough insulin-producing pancreatic islet cells remain that the diabetic disposition toward hyperglycemia does not yet exist.

The second notable process in patients with chronic disease is the pathophysiology in which the disease manifests. Benjamin Smart offers a "dispositional metaphysics of disease (DMD)" in which "diseases are causal processes best seen as simultaneously acting sequences of mutually manifesting dispositions" (2014, 252). Smart's account is similar to mine in terms of invoking dispositions. However, while I argue that for chronic diseases the disease is the disposition itself, for Smart the disease is the pathophysiological manifestation or "causal process" produced by dispositions. He uses the example of epilepsy, a chronic disease. On Smart's DMD, epilepsy is the seizure process that results when the body's disposition toward seizures and the disposition of some stimulus to provoke seizures mutually manifest themselves. However, epileptic seizures are realized only some of the time. In between seizures, the patient still has epilepsy. Rather than the pathophysiologic process in which the bodily disposition toward seizures manifests, epilepsy is the bodily disposition itself.

Marc Lange also proposes a sort of dispositional account of disease: "diseases as natural kinds of incapacities" (2007, 278). For Lange, an incapacity is a *lack* of a particular disposition; for instance, "PKU [phenylketonuria] is the incapacity to make enough active pheOH" (phenylalanine hydroxylase; 2007, 276), the enzyme that normally converts phenylalanine to tyrosine. It is questionable whether an incapacity, as the absence of something, is an entity. If we permit the existence of absences as entities, then we populate the universe with an infinite number of things: the absence of unicorns, the absence of temperatures below zero degrees Kelvin, and so on. To avoid this metaphysical baggage, we can often construe an incapacity as a capacity. We can redefine PKU as "the capacity or disposition toward insufficient active pheOH." Systematically implementing this adjustment to Lange's account would bring it in line with my own.[21] It would also cohere with the common medical practice of defining diseases through their positive harms (e.g., airflow limitation in COPD, hyperglycemia in diabetes mellitus).

Finally, Neil Williams argues that "some diseases are best thought of as sets of dispositions, whereas others are best thought of as processes" (2011, 211). On Williams's account, the body's cells form causally interacting networks and normally possess a set of typically manifesting "homeostatic

[21] Richard Scheuermann and colleagues (2009) propose that all diseases are dispositions. However, their objective is to provide a new standardized disease taxonomy/ontology, while my objective is to describe the current nature of chronic diseases.

dispositions" along with dispositions that are not typically manifesting. In typical environments, the typically manifesting dispositions produce typical processes. However, in atypical environments, the dispositions that do not typically manifest produce atypical processes. These atypical processes are diseases. Williams provides the example of scurvy, in which lack of vitamin C in the cellular environment disrupts cellular synthesis of collagen, leading to manifestations of a lack of collagen such as the breakdown of connective tissue and bleeding gums (which, for Williams, is scurvy). Other diseases, such as PKU, consist either in the addition of dispositions that are abnormal for the cell to have or the lack of normal dispositions. So, in total, for Williams there are three kinds of diseases: abnormal dispositions, the absence of normal dispositions, and processes resulting from normally inactive dispositions.

Williams suggests that the absence of a normal disposition toward some manifestation can be reconstrued as the presence of an abnormal disposition toward a different manifestation, so his second kind of disease is covered by his first kind. Furthermore, his suggestion that some *acute* diseases are processes is perfectly consistent with my theory of *chronic* diseases. What makes processes like the one in scurvy acute is that they are associated with an atypical and remediable condition of the environment (e.g., lack of dietary vitamin C). In contrast, what makes dispositions like PKU chronic is that they result from an irreversible bodily condition (e.g., genetic inheritance).

We have concluded our tour of rival theories. Having ruled out the plausible alternatives, we must again conclude that chronic diseases are generally dispositions.

## 4.4. Mental Disorders Are Dispositions

Mental disorders generally resemble chronic diseases in their chronic time course and are equally a fixture of the new modern medicine. Mental disorder diagnoses such depression are common worldwide (though skeptics have argued that some of the diagnosed depression results from overdiagnosis (Horwitz and Wakefield 2012)). I will now examine whether the dispositional theory is a good fit for mental disorders, and a better fit than the syndrome and fictionalist alternatives.

In probing the nature of chronic diseases, we can gather evidence from a variety of medical textbooks, articles, and practice guidelines, but there

is no authoritative diagnostic resource for all diseases or chronic diseases. In contrast, in scientific medicine there are two authoritative diagnostic resources for mental disorders: the *Diagnostic and Statistical Manual of Mental Disorders* (DSM), which is produced by the American Psychiatric Association (APA) and is especially authoritative in North America, and the *International Classification of Diseases* (ICD), produced by the WHO. (While the ICD also includes diseases, it is rarely used for clinical diagnosis outside of psychiatry.) As the DSM has been more commonly studied by philosophers of psychiatry, I will examine mental disorders as they are characterized in the DSM, though I would likely reach the same conclusions using the ICD due to important diagnostic similarities between the two resources and the attempts by DSM and ICD working groups to harmonize the manuals (APA 2022).

The DSM was first published in 1952 and is now in its fifth edition, DSM-5 (APA 2013), with a text revision (TR) to the fifth edition published in 2022, DSM-5-TR (APA 2022). While the first two editions of DSM were heavily influenced by psychoanalytic psychiatry, the dominant North American tradition in the early post–World War II era, the DSM-III, published in 1980, resulted from a dramatic shift in diagnostic approach. In recognition that the pathology of a mental disorder was often unknown or speculative, the designers of the DSM-III focused their efforts on establishing reliable diagnostic criteria for each disorder. The diagnostic structure of DSM-III was then replicated by all subsequent editions of DSM produced to date. The DSM's new approach to psychiatric classification is descriptive and atheoretical in that it relies on observable features of mental disorders that regularly cluster together, and it is generally agnostic as to their causes (Tsou 2011). Rather than their etiology or pathology, mental disorders are generally diagnosed using diagnostic criteria describing the presence or absence of certain symptoms and behaviors (though sometimes a cause of the disorder is also among the criteria, as in *posttraumatic* stress disorder, *etiological subtypes* of major neurocognitive disorder, and *substance-induced* mental disorders).

For example, the nine DSM-5 criteria for depression are depressed mood, markedly diminished interest or pleasure in activities, significant unintentional weight loss/gain or decreased/increased appetite, insomnia or hypersomnia, psychomotor agitation or retardation, fatigue or loss of energy, feelings of worthlessness or inappropriate guilt, difficulty concentrating or indecisiveness, and suicidal ideation or suicidal behavior (APA 2022). The

first eight symptoms/behaviors must be present most of the time. When at least five of these criteria, including at least one of the first two, are satisfied over a period lasting at least two weeks (and certain other criteria are satisfied),[22] the patient meets the criteria for a "major depressive episode" and is diagnosed with depression.

According to DSM-5, "Diagnostic criteria are offered as guidelines for making diagnoses" (APA 2013, 21). However, the manual also states that "the DSM criteria define disorders" (xii), or at least "summarize characteristic syndromes" (19). In other words, the criteria not only are diagnostic criteria but also tell us what makes a certain disorder a case of depression versus, say, a case of a substance use disorder. While psychiatrists and philosophers have criticized the treatment of DSM criteria as definitive (Hyman 2010; Zachar 2014; Kendler 2017), the criteria are nonetheless used in DSM and by many of its users to characterize mental disorders and not only to diagnose them. Thus the DSM and its criteria offer a promising glimpse into the nature of mental disorders.

In its introductory section, the DSM-5 authors comment on the purpose of the manual and how it is to be used. Here we find conflicting statements about the nature of mental disorders. In their definition of a mental disorder, the authors write, "A mental disorder is a *syndrome* characterized by clinically significant disturbance in an individual's cognition, emotion regulation, or behavior that reflects a dysfunction in the psychological, biological, or developmental processes underlying mental functioning" (APA 2013, 20, my emphasis). Presumably, these syndromes consist of the symptoms and behaviors listed in the diagnostic criteria for each disorder. However, in other places the authors write as if mental disorders are the cause of these syndromes. For instance, they say that the DSM criteria "are intended to summarize characteristic syndromes of signs and symptoms that point to an underlying disorder" (19) and that the validity of a diagnosis refers to "the degree to which diagnostic criteria reflect the comprehensive manifestation of an underlying psychopathological disorder" (5).

[22] Additional criteria establish that: the symptoms cause significant distress or impairment in functioning; they are not attributable to a substance or medical condition; they are not better explained by another mental disorder; and there has never been a manic or hypomanic episode (APA 2022). The first three of these additional criteria are standardly found among the DSM criteria for a mental disorder.

Because syndromes are central to DSM diagnosis, these two broad views—the "syndrome view" and the "syndrome-cause view"[23]—are the most obvious starting points for a theory of mental disorders (Maung 2016; Kendler 2017). The syndrome view holds that mental disorders are syndromes, while the syndrome-cause view holds that mental disorders are the cause of these syndromes. Each of these alternatives admits several more fine-grained possibilities. On the syndrome view, the symptoms and behaviors constituting the mental disorder could co-occur because they have a common cause (an underlying psychopathology) or because they cause one another (they form a "causal symptom network"). On the syndrome-cause view, the mental disorder that causes these symptoms could be any one of the entities considered in the previous section (e.g., a process, a property).[24] I will first compare these two broad views and will discard the syndrome view, before endorsing a more fine-grained theory of mental disorders from under the umbrella of the syndrome-cause view. Evidence against the syndrome view and in favor of the syndrome-cause view is provided by the way the DSM authors intend for their diagnostic criteria to be applied and, in particular, how they accommodate silent disorders.

Earlier I introduced the nine criteria that define a major depressive episode, a variable syndrome consisting in depressed mood, diminished interest in activities, and other manifestations. The characteristic syndrome of type 1 bipolar disorder is called a manic episode and consists of features such as an abnormally elevated, expansive, or irritable mood and increased energy or goal-directed activity. Meanwhile, the characteristic syndrome of type 2 bipolar disorder is an alternating sequence of major depressive episodes and hypomanic episodes (the less severe equivalent of a manic episode). The characteristic syndrome of other mental disorders is not always episodic and does not always have a name. For instance, the characteristic syndrome of generalized anxiety disorder is excessive and uncontrollable

[23] Jennifer Radden (2003) refers to the syndrome view as "ontological descriptivism." These views are related to two of the three major metaphysical theories in the psychiatric literature discussed by Peter Zachar (2019), which he calls "the causal network model," in which the "real clinical entity" is a causal symptom network, and "the medical model," in which the "real clinical entity" is the underlying pathology responsible for the symptoms. The third major theory discussed by Zachar is the "factor analytic model," in which the "real clinical entities" are dimensional variables. I will set aside the factor analytic model because—as the DSM-5 authors repeatedly emphasize (APA 2022)—DSM mental disorder diagnoses are categorical rather than dimensional.

[24] The syndrome-cause view could also allow that the symptoms/behaviors are part of a causal network. It need only commit to the view that among the nodes in this network is a cause that isn't one of the symptoms or behaviors listed in DSM diagnostic criteria, and that this node is itself the mental disorder.

anxiety and worry on most days regarding several events or activities (for at least six months), along with at least three of six other manifestations. Table 4.2 includes further examples.

The DSM implicitly distinguishes between a mental disorder and its characteristic syndrome, as evidenced by the manual's own silent disorder phenomenon. To receive an initial diagnosis of a mental disorder, its characteristic syndrome must initially be present. But the diagnosis can outlast its syndrome. Many mental disorders are recognized by DSM-5 as having a remitting-relapsing course, or at least having periods of partial remission, perhaps due to treatment. The manual even includes diagnostic specifiers for patients with depression in "full remission" or "partial remission" (e.g., "major depressive disorder in full remission"). A full or partial remission is one in which all or some of the symptoms/behaviors associated with the disorder disappear, at least temporarily. Yet the diagnosis of depression remains glued to the patient under the implicit assumption that they have not recovered from depression but are merely experiencing what may prove to be a temporary relief from a depressive episode. This practice is coherent if depression is the cause of a depressive episode (a cause that may go silent) but not if depression is a depressive episode.

Further support for the syndrome-cause view over the syndrome view is provided by the manner in which DSM-5 handles the diagnosis of bipolar disorder. In order to distinguish depression from bipolar disorder, the criteria for depression exclude a history of manic or hypomanic episodes. If a patient has had a depressive episode and has never had a manic or hypomanic episode, they may receive a diagnosis of depression. If they have had a depressive episode as well as a manic or hypomanic episode, they instead receive a diagnosis of bipolar disorder (type 1 or type 2, respectively). However, sometimes a first depressive episode will occur before a first manic or hypomanic episode. Thus a patient could initially meet the criteria for depression and receive a diagnosis of depression, but later have their diagnosis revised to bipolar disorder (Nestsiarovich et al. 2021). DSM-5 describes this diagnostic switch as follows: "Many bipolar illnesses begin with one or more depressive episodes, and a substantial proportion of individuals who initially appear to have major depressive disorder will prove, in time, to instead have a bipolar disorder" (APA 2013, 165). According to the DSM, it is not that the patient had depression and then developed a bipolar disorder, as one might say if depression and bipolar disorder were identical to their characteristic episodes. Rather, the patient was initially misdiagnosed (APA 2022) and had

**Table 4.2** Examples of Dispositional Mental Disorders

| Mental Disorder | Characteristic Manifestation |
|---|---|
| Major Depressive Disorder | Depressed mood, anhedonia, increase/decrease in weight/appetite, insomnia/hypersomnia, psychomotor agitation/retardation, fatigue/loss of energy, inappropriate guilt or feelings of worthlessness, decreased concentration or indecisiveness, suicidality |
| Generalized Anxiety Disorder | Excessive and uncontrollable anxiety and worry about a number of events or activities, restlessness or feeling on edge, easily fatigued, difficulty concentrating or mind going blank, irritability, muscle tension, sleep disturbance |
| Bipolar 1 Disorder | *Manic episode*: abnormally and persistently elevated/expansive/irritable mood and persistently increased goal-directed activity/energy, increased self-esteem or grandiosity, decreased need for sleep, talkativeness or pressured speech, flight of ideas, distractibility, increase in goal-directed activity or psychomotor agitation, excessive involvement in activities with painful consequences (symptoms result in significant dysfunction or hospitalization, or have psychotic features) |
| Bipolar 2 Disorder | *Major depressive episode* and *hypomanic episode* (same symptoms as manic episode but without significant dysfunction and hospitalization and without psychotic features) |
| Schizophrenia | Delusions, hallucinations, disorganized speech, disorganized or catatonic behavior, negative symptoms (flat affect or avolition) |
| Major Neurocognitive Disorder (Dementia) | Diminished performance in at least one cognitive domain (executive function, complex attention, language, learning and memory, social cognition, perceptual-motor) |
| Substance Use Disorder | Using the substance more than intended, failed attempts to cut back on use, significant investment of time related to substance use, cravings, failure to fulfill major obligations, continued use despite social consequences, important activities given up or reduced, use of substance in dangerous situations, use continues despite knowledge of its harms, tolerance, withdrawal |
| Obsessive-Compulsive Disorder | Obsessions, compulsions |
| Autism Spectrum Disorder | Deficits in social communication and interaction (social-emotional reciprocity, nonverbal communication, relationships), restricted and repetitive interests or activities (stereotypic or repetitive behavior, rigid routines or ritualized behavior, restricted and fixated interests that are abnormal in intensity or focus, unusual sensory reactions or sensory engagement). |

*Source:* APA (2022)

bipolar disorder all along even though they *appeared* to have depression, suggesting that bipolar disorder is a cause of manic/hypomanic episodes (as well as depressive episodes).

These two observations favor the syndrome-cause view. Among the many possible kinds of things the syndrome's cause may be, the fact that mental disorders are characterized by their effects or "manifestation" (APA 2013, 5) lends support to the dispositional theory. Rather than an independently identifiable psychopathology, mental disorders are identified through their manifest syndromes. Because DSM criteria define (or at least characterize) their disorder and describe a syndrome, Hane Maung (2016) infers that DSM diagnoses must (at least sometimes) refer to syndromes. However, as we saw with chronic diseases, clinical findings can define a disorder without being that disorder when the disorder is a disposition and those findings are its characteristic manifestation.[25] Just as asthma can be understood as a disposition toward a state giving rise to an asthma attack and heart failure can be understood as a disposition toward signs and symptoms of cardiac congestion, depression is a disposition toward a major depressive episode, bipolar disorder is a disposition toward a manic/hypomanic episode, and generalized anxiety disorder is a disposition toward excessive and uncontrollable anxiety/worry.[26] Robert Kelly develops a unifying account of addiction as "*the systematic disposition to fail to control one's desires to engage in certain types of behaviors*" (2023, 22).

Because individuals diagnosed with a mental disorder often retain their diagnosis for a long duration, if not indefinitely, mental disorders are generally treated as chronic disorders. For some progressive disorders such as major neurocognitive disorder (hereafter referred to as "dementia"), patients

[25] Maung (2016) argues for a two-dimensional semantics in which psychiatric diagnoses sometimes refer to syndromes, other times to their underlying pathophysiological cause. However, the dispositional theory reconciles these two dimensions because a psychiatric disposition is a cause of a syndrome that is diagnosed by identifying that syndrome. Moreover, as Maung acknowledges, sometimes the pathophysiological cause of a psychiatric syndrome can be variable. In these cases, Maung suggests that the diagnosis refers to a disjunction of pathophysiological causes. Considering that what the disjuncts have in common is that they are disposed to manifest the same syndrome, this brings Maung's account a bit closer to my dispositional theory.

[26] The dispositional theory of mental disorders shares similarities with diathesis-stress models in the psychiatric literature (e.g., Monroe and Simons 1991). In diathesis-stress models, an individual vulnerable to mental illness has a "predisposition" or "diathesis" for that illness, often thought to be present from early life. The diathesis may later be triggered by life stress and result in mental illness. However, it is often unclear in these models (and is an empirical question) whether the diathesis is a mental disorder disposition or a disposition toward a mental disorder disposition. Sanja Dembić (2023) argues that mental disorders are a different sort of modal property, an *inability*—more specifically, the inability to respond adequately to available reasons, resulting in harm.

typically do not experience complete long-term remission (Emmady et al. 2022); for others, such as bipolar disorder, a relapsing-remitting untreated course is common (Winokur et al. 1993). In these cases, it is reasonable to assume that the disposition persists and that the disorder is truly chronic.[27] Other disorders, such as depression, have a more variable natural history, and it is possible that a patient's symptoms of depression may never return (B. Stegenga et al. 2012). Thus, when their symptoms resolve, it could be that their disorder has gone silent and persists unmanifested (they are in remission), but it is also possible that they have recovered from their disorder (Monroe and Harkness 2011). Lacking independent and reliable markers for the persistence of the disposition beyond its manifestations, it is difficult to distinguish among these possibilities.

One challenge for a dispositional view of DSM disorders (and also for a dispositional view of certain chronic diseases) is that DSM disorders are diagnosed only after their manifestations are present, yet a disposition can (and usually will) be present before it manifests. If we think that DSM disorders are dispositions, then the result is the delayed diagnosis and potentially also the underdiagnosis of mental disorders. There is delayed diagnosis because mental disorders are diagnosed only after they manifest, and potentially underdiagnosis because mental disorders that never manifest are never diagnosed. Is it plausible, then, that DSM criteria refer to dispositions? We do not, for comparison, consider an object to be fragile only after it breaks. The solution to this puzzle is that DSM criteria not only characterize a disorder but also (and foremost) serve as its diagnostic criteria, or rules for making a diagnosis. There are practical reasons for including the presence of symptoms and behaviors among DSM diagnostic criteria, especially the improvement of diagnostic interrater reliability in the absence of diagnostic tests that can latch onto an underlying psychopathology (Hyman 2010). Rather than casting doubt on the dispositional theory of mental disorders, the DSM practice of diagnosing disorders only after symptoms and behaviors are present reflects their function as diagnostic criteria. Any delayed diagnosis or underdiagnosis can be seen as a byproduct (and not necessarily a harmful one, given that psychiatry generally lacks means for identifying the

[27] Whether chronic mental disorders are chronic *diseases* is a question that my analysis cannot answer because an answer requires grappling with our concept of disease. However, by arguing that mental disorders are not syndromes, I have cleared the way for one important barrier to considering them diseases.

presence of mental disorders independently of their manifestations and of preventing their progression from subclinical disorders to clinically manifest ones).

A peculiarity of the DSM diagnostic criteria complicates its disposition-based classification system. The symptom and behavior criteria for a DSM diagnosis are often polythetic: not all of the criteria are necessary, and satisfying a given number of them will satisfy the criteria for diagnosis (APA 2022). For example, the criteria for depression are met when at least five of nine criteria are satisfied (including at least one of the first two criteria). One patient might satisfy criteria 1 through 5, while a second might satisfy criterion 2 plus criteria 6 through 9 (not to mention any further differences in symptoms beyond those listed in the criteria).

We could define two subtypes of depression corresponding to these two subtypes of disposition, and then the type "depression" refers to a family of disposition subtypes. These subtypes of depressed disposition might depend on different kinds of causal base. Or they might depend on the same kind of causal base that manifests differently in different individuals due, for instance, to its interaction with other causes that are present in some cases but not in others. Whether these different dispositions are rightly considered subtypes of what is fundamentally the same type of disposition—depression—might depend on whether they share the same causal base. It might also depend on how we characterize the depressed phenotype(s), including how abstractly. A disposition toward sadness and a disposition toward anhedonia might be considered the same type of disposition—depression—if we characterize low mood and anhedonia more abstractly as the same manifestation, "depressed mood." Analogously, a disposition to crack and a disposition to shatter can be considered the same type of disposition—fragility—if we consider cracking and shattering to be variations on the same abstract manifestation, breaking.

For many mental disorders, their causal base or "psychopathology" is poorly understood. For others, it is comparatively better understood. For example, there are multiple subtypes of dementia, each depending on neurodegeneration in a different part of the brain or resulting from a different etiology. Thus different subtypes of dementia can have a different psychopathological causal base: the psychopathology of Alzheimer's disease includes widespread cortical atrophy, while the psychopathology of frontotemporal degeneration includes degeneration of the frontal and temporal lobes (Emmady et al. 2022). In contrast, the psychopathology of

specific phobia is dysregulation in the mind's fear system due to an improperly learned stimulus-response association (Pereira 2020). The dispositional theory of mental disorders is not committed to any particular science or "level" of scientific description, nor to any particular psychiatric research program such as biological psychiatry (Murphy 2011). In principle, a disposition toward a cluster of symptoms and behaviors could depend on a molecular, neurological, or cognitive psychopathology.

The metaphysical alternatives to the dispositional theory canvassed (and rejected) in the previous section in my discussion of chronic diseases encounter similar problems in the context of mental disorders, and I will not repeat those problems here. However, fictionalism about mental disorders demands special consideration. While there is certainly skepticism about the legitimacy of some chronic diseases discussed in the previous section, such as borderline cases of hypertension, type 2 diabetes, and osteoporosis (Schwartz 2008), that skepticism is isolated to some instances of some chronic diseases and cannot threaten the validity of my generalization about the nature of chronic diseases. However, skepticism about the legitimacy or "reality" of psychiatric diagnoses is more common and sometimes more thoroughgoing, from the "antipsychiatry" movement of the 1960s and 1970s to the "neurodiversity" and "mad pride" movements of today. While I cannot do justice to the diverse range of critical views here, one recurring theme among skeptical critiques of psychiatric diagnoses is that they do not pick out genuine disorders but merely something like socially unexpected behavior or divergent mental functioning. For example, in famously declaring mental illness to be a "myth," Thomas Szasz (1961) was partly claiming that patients treated by psychiatrists do not always have an "illness" (which Szasz took to mean a physical lesion or a biological dysfunction).[28]

If there are *no* mental disorders, then the claim that mental disorders are dispositions cannot be true (at least, if we interpret that claim as having existential import). However, most of psychiatry's critics are selective skeptics rather than blanket skeptics: they do not doubt all psychiatric diagnoses but rather argue that particular diagnoses are not mental disorders. For instance, neurodiversity proponents argue that certain divergent cognitive states, like

[28] Szasz (1961) also argued that the conditions treated by psychiatrists cannot be illnesses because they are behaviors, and illnesses are not behaviors. Szasz's argument relies on the syndrome view of mental disorders, which I have argued faces serious difficulties.

(some cases of ) autism, shouldn't be considered mental disorders (Chapman 2019); Allan Horwitz and Jerome Wakefield (2012) argued that some cases of sadness diagnosed as depression are not mental disorders; and in the 1970s, R. D. Laing and A. Esterson (1970) argued that people diagnosed with schizophrenia had nothing wrong with them. Even Szasz accepted that at least some conditions treated by psychiatrists are in fact illnesses, such as the brain lesions of neurosyphilis (syphilis affecting the brain). If we grant that there are at least some mental disorders, then fictionalism does not undermine the claim that mental disorders are dispositions. Those DSM diagnoses that refer to genuine mental disorders refer to dispositions, insofar as every DSM category has a structure similar to that of the examples that I used to motivate the dispositional theory.

Even if we accepted that there are no mental disorders, the main question I have been investigating is not really about conditions comprising the category "mental disorder" but rather about the referents of DSM diagnoses: "To what do psychiatric diagnoses refer?" (Maung 2016). My main conclusion is that these diagnoses refer to dispositions. When I say that mental disorders are dispositions, I really mean that the entities diagnosed according to the *Diagnostic and Statistical Manual of Mental Disorders* are dispositions, regardless of whether the DSM has been appropriately labeled as a manual of mental disorders or whether instead its diagnoses refer to non-disordered psychological dispositions.

However, this way of rephrasing the question and my answer will not evade one line of fictionalist argument that goes as follows. What I have shown through my analysis is only that the DSM authors assume there are dispositions responsible for DSM syndromes. I have not shown that those syndromes in fact result from dispositions. Dispositions—or at least the medical dispositions in which I am interested—are not just any property. They are properties that realize their manifestations only under certain stimulus conditions. The DSM is not fully explicit about what these stimulus conditions are supposed to be. But the manual often excludes substances, medical conditions, and common stressors or losses. When we say that an object is fragile, or disposed to break under stress, we mean that the object will break under certain kinds or amounts of stress that will not cause objects lacking this disposition to break. Likewise, when I claim that people receiving a DSM diagnosis of depression are disposed toward a depressive episode, I mean that they manifest a depressive episode given stressors (or even in the absence of stressors) that will not cause people lacking this

disposition to have a depressive episode. To say that an object is fragile (or a person has depression) is *not* to say that the object will break (or the person will experience depressive symptoms) under just any circumstances, for any object can break (and any person can metaphorically break) under extreme stress.[29]

It could be argued that the symptoms and behaviors described in the DSM do not result from unusual dispositions possessed only by those meeting DSM diagnostic criteria and manifesting in the absence of common stressors, but rather from stressors that could cause someone without an abnormal disposition to experience those same symptoms and behaviors. Arguments roughly along these lines were made by Laing and Esterson (1970) about behaviors diagnosed as schizophrenia, and by Horwitz and Wakefield (2012) about some cases of grief diagnosed as depression under DSM-IV. Horwitz and Wakefield's diagnosis of the problem is that these cases of grief are not properly diagnosed as depression because they are not mental disorders. Again, my central question is not really whether DSM diagnoses refer to mental disorders but whether they refer to dispositions. Nonetheless, their argument is still relevant for the dispositional theory even if their conclusion is not. If grief-stricken individuals diagnosed with depression are reacting to common stressors rather than manifesting unusual dispositions, then DSM criteria for depression may misattribute a disposition to them that doesn't exist. Diagnosing someone with a disposition upon finding symptoms that such dispositions manifest is a fallible inference, akin to diagnosing an object as fragile after finding its broken pieces.

In response to this objection, I will qualify my thesis to allow that DSM criteria (or those people applying them) might sometimes misattribute a disposition to individuals diagnosed. DSM diagnoses *generally* refer to dispositions. When (and thus how often) misattribution occurs is difficult to determine absent reliable markers for mental disorders and their psychopathology as well as a better understanding of the causes of psychiatric syndromes.

[29] Robert Kelly (2023) develops this idea for the addiction disposition using the concept of "systematicity." Addiction is not simply failing to control one's desire for an addictive substance or behavior. Addiction is a systematic disposition toward failure to control one's desires, that is, a disposition that reliably manifests in enough stimulus conditions that people with the disorder are highly likely to find themselves in. (For addiction, Kelly says these stimulus conditions are societally relevant "choice situations.")

## 4.5. Dispositions and Specific Pathology

The dispositional nature of chronic diseases and mental disorders can create challenges for their treatment and prevention. In particular, dispositions may violate the principle of specific pathology that for other diseases has facilitated the development and deployment of targeted medical interventions.

Effective prognosis, treatment, and prevention of disease often hinges on making inductions from past cases of disease, especially past cases followed in research studies. (It also depends on having good prognostic markers and effective medical interventions.) When inductions are made from past instances of a disease type to other instances of that type and the inductions are sound, then those disease categories can be considered "projectible" with respect to these inductive inferences (N. Goodman 1954). Projectability is a virtue often thought to be possessed by natural kinds. For instance, Williams (2011) argues that our disease categories are projectible and thus are natural kinds, as evidenced by the fact that inferences about the success of treatment in a new case of disease based on prior success with that treatment in another case of the disease are often accurate. (See also Tsou 2021, 2022 on the projectability of some psychiatric kinds.) On Williams's account, though multiple properties determine membership in a disease natural kind,[30] the most important properties for determining disease kind membership are bodily dispositions (see also Lange 2007).

Similarly, Rachel Cooper (2014, 63) argues that while many mental disorders are not well enough understood to determine whether they are natural kinds, "[s]ome types of mental disorder are plausibly natural kinds." For Cooper, it would be beneficial if mental disorders were natural kinds because then we could use them to make explanations and predictions in psychiatry:

> If bipolar disorder, say, turns out to be a natural kind, then we can explain why it is that a particular case unfolds as it does. We can say that someone suffers from alternating manic and depressive phases *because* they have bipolar disorder. We will also be able to make various predictions, for example, if someone suffers from bipolar disorder then lithium treatment might help. Where mental disorders are natural kinds, this will give us some control over them. In particular, treatments that work for one of the kind can be expected to work for all of the kind. (63–64)

[30] Williams argues that disease natural kinds are homeostatic property cluster kinds, drawing on the work of Richard Boyd (1991).

Cooper considers an objection to the thesis that mental disorders are natural kinds that she calls "the functionalist's objection," adapted from McGinn (1991). The functionalist's objection and Cooper's response are each part of broader arguments. I will rehearse only those parts that are relevant to the discussion in this section.

Colin McGinn argues that the functionalist cannot maintain that mental kinds are natural kinds. Functionalism about the mind is the view that mental states are characterized by their functional role: their sensory inputs, behavioral outputs, and interactions with other mental states. (For the functionalist, depression might be the state that leads some people to stay in bed all day regardless of stimuli that might motivate others to go to work in the context of other mental states representing self-deprecating beliefs, etc.) McGinn argues that according to functionalism, the same mental state is multiply realizable: it can conceivably be realized by multiple systems (other living systems, other nonliving systems). A mental state could thus be realized by very different structural properties in different instances even while performing the same functional role. On the other hand, our prototypical examples of natural kinds (e.g., chemical elements) don't vary with respect to the structural properties of their realizers because they are grouped together according to their structural properties (e.g., atomic number). Therefore, argues McGinn, multiple realizability gives the functionalist good reason to reject the view that mental states are natural kinds.

Cooper's response is basically the idea that multiple realizability does not entail multiple realization. Even if a mental state like pain could in principle be realized differently in nonhuman systems, it may in fact be realized similarly in all humans. Likewise, a type of mental disorder could be realized by the same brain abnormality in all humans. If depression were always realized by an abnormality in neurotransmitter levels, it would not be multiply realized, Cooper suggests. Cooper acknowledges that it might turn out that certain mental disorders such as schizophrenia do have distinct kinds of realizers. In these cases, the mental disorder type would cover multiple natural kinds. However, the functionalist's objection does not rule out the possibility that some of our current kinds of mental disorders always depend on the same kind of internal structure and thus are typical natural kinds.[31]

[31] However, McGinn adheres to Saul Kripke's (1972) account of natural kinds and argues that the mere possibility of multiple realization of mental kinds (i.e., that they are conceivably multiply realizable) is a good a priori argument against their being natural kinds because natural kinds *necessarily* are not multiply realizable.

While McGinn and Cooper discuss multiple realizability as a potential problem for functionalism about the mind, there is a clear analogy with dispositions. As discussed in section 4.3.1, disease dispositions depend on a pathological causal base. Token instances of a disease disposition may even be identical with their token causal base. And yet a type of disease may not be identical to a type of causal base because dispositions may be multiply realized (Prior et al. 1982). Different instances of the same type of disposition may depend on different types of causal base. One object may be breakable because it is porous, while another may be breakable because it is thin. In fact, multiple realization may be a more serious worry for disease dispositions compared to functionalist mental states because while functionalist mental states are characterized by their causes and effects, disease dispositions are conventionally characterized only by their effects, and thus the range of potential causal bases on which a dispositional disease might depend is less constrained.

Williams and Cooper are interested in whether diseases or mental disorders are natural kinds at least partly because disease natural kinds are thought to be projectible for the purposes of explanation and prediction. As I will argue in a moment, concerns about multiple realization or multiple causal dependency are indeed relevant for these kinds of medical inferences. Rather than a problem for the thesis that diseases are natural kinds, we can thus treat this problem more directly as a problem for explanatory and predictive inferences in medicine. Unfortunately, while Cooper is right that the functionalist's objection does not rule out the possibility that our current mental disorders (and diseases) always depend on the same kind of categorical structure, her response does not give us any reason to think that many mental disorders (and diseases) *do in fact* always depend on the same kind of categorical structure.[32]

Rather, we should worry that chronic diseases and mental disorders commonly depend on a variable causal base or bodily pathology. Just as two or more distinct causal bases can ground fragility on different occasions, two or more bodily pathologies can ground chronic medical conditions on different occasions. The same physiologic manifestation can often be produced by different pathologies. The problem is magnified when those manifestations are

[32] Cooper mentions Huntington's disease as one plausible example of a mental disorder that is a natural kind because it depends in all cases on the same defective gene. In DSM-5, Huntington's disease is listed as an etiological subtype of major neurocognitive disorder (APA 2022).

nonspecific symptoms, including somatic symptoms like shortness of breath as well as the symptoms of mental disorders. Thomas Insel, the former director of the National Institute of Mental Health of the National Institutes of Health in America, encapsulates the problem: "[D]iagnosis based only on symptoms may never yield the kind of specificity that we have begun to expect in the rest of medicine. Behavioral symptoms are multidetermined, so diagnoses based only on presenting complaints are unavoidably heterogeneous in terms of pathophysiology" (2014, 395). Moreover, because the criteria for diagnosis of mental disorders are generally polythetic, two individuals meeting the criteria for the same DSM diagnosis may differ sharply in their symptoms and presentation. This could further magnify the problem of variable psychopathology underlying a given DSM diagnosis (Tabb 2015).

Insel may exaggerate in saying that symptom-based classification *unavoidably* results in variable pathology or pathophysiology. Some symptoms or symptom clusters could in principle be caused by one or only a few pathologies. For instance, neurological syndromes often result from lesions in specific areas of the nervous system. (Parkinsonism is most often due to degeneration in the basal ganglia; Shrimanker et al. 2022.)

Unfortunately, for many chronic diseases it is known that different instances of the disease depend on different pathology, and sometimes disease subtypes are even defined according to a distinct pathology. Consider COPD, which includes three subtypes: emphysema, chronic bronchitis, and small airways disease. Each of these subtypes of COPD depends on a different type of lung pathology. Emphysema depends on destruction and enlargement of lung alveoli, while small airways disease depends on narrowed small bronchioles. Similarly, diabetes mellitus has many subtypes, including type 1 and type 2. The causal base of type 1 diabetes is deficiency of insulin-secreting cells, while type 2 diabetes depends on dysfunction in insulin-responsive cells and sometimes also in insulin-secreting cells. Heart failure sometimes depends on a cardiomyopathy (muscular pathology), and sometimes on a stenosis or narrowing of the aortic valve. Finally, dementia, a mental disorder, comes in several varieties with distinct brain pathology, including Alzheimer's disease (in which there is widespread cortical atrophy) and frontotemporal dementia (in which the frontal and temporal lobes degenerate).

For some types or subtypes of disease, the pathology is poorly understood, and for all we know it could be highly variable. Examples include

type 2 diabetes, essential hypertension, heart failure, and the majority of mental disorders. Pathological variation may lurk in the fog of our poor understanding and may even explain why we lack understanding because it is difficult to understand a disease or disorder as a singular kind when its pathological basis is not singular. Even for disease types or subtypes with some shared pathology, not enough pathological features may be shared for prognosis, treatment, or prevention to be reliable. In other words, there may remain important differences in pathology even among cases with similarities significant enough to characterize them as having shared "specific pathology."

In Chapter 3, I described a constitutive model of classification in which disease kinds are distinguished according to the bodily condition of which the disease consists. For infectious diseases and cancers, this condition is typically a specific pathology, a singular and necessary categorical bodily state that all cases of that disease type share. Infectious diseases are classified as an infection with a specific pathogen, while cancers are classified as tumors or proliferations of a particular cell and tissue type (histopathology). I argued that classification by specific pathology can help with the development and deployment of effective targeted medical interventions such as pathogen-specific vaccines and targeted cancer treatments. A targeted treatment interacts with the disease in a pathology-specific manner the way that vaccine-generated antibodies interact with a specific antigen on a pathogen or the way that monoclonal antibodies interact with a specific antigen on a cancer cell. Chronic diseases and mental disorders are also classified using the constitutive model; they are classified as a particular bodily disposition. In their case, the resulting *lack* of specific pathology can hinder the development and deployment of effective interventions.

When a disease or disorder lacks specific pathology, medical interventions meant to cure or relieve its symptoms may have highly variable effectiveness, highly heterogeneous treatment effects. That is because these interventions interact with the disposition's causal base. Surgeries interact with the anatomical structure underlying a disease, while medications interact with the underlying pathophysiology. For example, angioplasty, a treatment for coronary artery disease, interacts with narrowed coronary arteries, while selective serotonin re-uptake inhibitors, a treatment for anxiety and depression, interact with the serotonin system in the brain. Cooper suggests that when a disease is a natural kind, "treatments that work for one of the kind can be expected to work for all of the kind" (2014, 64). The complete absence of variation in treatment effectiveness suggested by this quote may be too

demanding; it may require so little pathological variation among patients with the same diagnosis that diagnostic categories end up being practically bespoke, applying to only a small number of patients. However, if the causal base underlying a disease disposition is *highly* variable, the effectiveness of treatments that interact with the causal base might also be highly variable, which is a problem we should take seriously. Moreover, in research it may be more difficult to develop an effective targeted intervention, one that interacts with a specific pathology, when the disease lacks a specific pathology.

In Chapter 3, I argued that specific pathology was the key conceptual innovation that led to the discovery of universal etiologies for infectious diseases, and that universal etiology often provides a general causal explanation for the occurrence of disease and facilitates more effective preventive interventions. While specific pathology is no guarantee of universal etiology, nonspecific pathology is more likely to have variable etiology because, simply put, different bodily dysfunctions often have different causes. When a disease or disorder lacks specific pathology, preventive interventions might then be met with highly variable success because a preventive intervention disrupts an etiological mechanism (Chapter 6), and a targeted preventive intervention will disrupt a particular etiologic mechanism the way a vaccine prevents infection with a specific pathogen by generating antibodies against that pathogen. Moreover, studying a specific pathology that has a universal etiology can lead to the development of a more effective preventive intervention, one that disrupts an etiological mechanism common to cases of that pathology.

Finally, when a disease or disorder lacks specific pathology, its prognosis or symptoms might be highly variable because prognosis and symptoms also commonly depend on the pathology. For instance, different symptoms can occur in heart failure due to dilated cardiomyopathy versus hypertrophic cardiomyopathy, in COPD due to emphysema versus chronic bronchitis, and in dementia due to Alzheimer's disease versus frontotemporal dementia (where all of these conditions are associated with different kinds of disease pathology). Similarly, while hyperglycemia in type 1 diabetes is often severe and non-progressive (after the destruction of insulin-producing pancreatic cells ceases), hyperglycemia in type 2 diabetes is often initially milder and progresses over time.

Variability in treatment effectiveness and in prognosis is all too common in the care of patients with chronic disease or mental disorder (Kravitz et al. 2004). For example, different population groups—notably, men and

women—respond differently to carotid endarterectomy, the surgical treatment of atherosclerotic narrowing in a carotid artery (Rothwell et al. 2005). As another example, different patients with depression or anxiety may respond differently to antidepressants: one patient may experience improvement on one brand of antidepressant, a second patient with the same diagnosis may experience improvement on another brand (Connolly and Thase 2011). Variability in treatment success or in prognosis may be evidence of variation in the pathology underlying the disease.[33]

When patients with the same diagnosis share a common pathology, inferences about one patient from experience with others may be sound. However, "rational treatment"—the development of medical interventions through an understanding of disease—requires more: it requires that this common pathology is known in sufficient detail. Identifying a disease through its observable or measurable manifestations can be a way to tie diagnosis to the harms that give the disease its clinical significance. It can also be a way to enhance the reliability of diagnosis through readily observable findings and standardized criteria that include signs/symptoms and diagnostic test results. However, it may *also* sometimes betray a lack of understanding of the underlying pathology because, absent knowledge of the pathology, a disease's clinical features are often all we can use to identify the disease. (This is the explicit justification for the DSM's symptom-based classification system; Tsou 2011). Therefore the dispositional nature of a disease may portend difficulties for its rational treatment. (Of course, having detailed knowledge of disease pathology is no guarantee that medicine will succeed in developing effective treatment.)

When its pathological base is unknown, a disposition may also fail to offer a fully satisfying explanation for its manifestations. In the passage quoted earlier, Cooper suggests that when a disease is a natural kind, we can explain why a case unfolds as it does by invoking the disease: "We can say that someone suffers from alternating manic and depressive phases *because* they have bipolar disorder" (2014, 63). However, if all we know is that (type 1) bipolar disorder is a disposition toward manic and depressive episodes, saying that someone has bipolar disorder does not provide a fully satisfying

[33] In the example of carotid endarterectomy, variation in the morphology of the atherosclerotic plaque correlates with sex differences in treatment effectiveness (Rothwell et al. 2005). In general, causally relevant differences beyond the disease pathology can also explain differences in disease outcome. Relevant differences can include variation in healthy physiology, in the presence of concurrent diseases, in the use of other medical interventions, and in the social or physical environment.

explanation for why they have such episodes because it amounts to saying that they are disposed to having them. Granted, this explanation may not be completely unsatisfying because there is content to the claim that someone has a dispositional mental disorder that goes beyond saying "there was some cause of the depressive and manic symptoms present." As a disposition, a mental disorder is a stable property that depends on an enduring mental pathology. Saying that bipolar disorder was responsible for manic or depressive symptoms thus rules out other potential causes of these symptoms that don't depend on an enduring mental pathology, which include stimulant and depressive drugs. However, the explanation is not *fully* satisfying when the mental pathology is unknown, if we assume that to offer a satisfying causal explanation for an effect, the identity or nature of the cause must be known to us.

Overall, the dispositional nature of chronic diseases and mental disorders may spell trouble for medical explanation and prediction. Recall that chronic diseases are chronic partly because medicine lacks cures for them. Patients receiving a diagnosis of a mental disorder also typically retain that diagnosis indefinitely, implying that they too have a chronic disposition. If their dispositional nature figures in the reasons why chronic diseases and mental disorders are rarely cured (because the dispositions may lack known and specific pathology), then the dispositional theory also has something to say about why they are chronic. Consequently, the search for specific pathology that we can then target with medical intervention is a promising strategy for countering the contemporary epidemic of chronicity in medicine.

Earlier I mentioned the role of epidemiology in investigating the rising prevalence of chronic diseases, in facilitating the recognition and study of their chronicity, and in establishing asymptomatic risk factors as chronic diseases. The shift away from laboratory medicine and toward epidemiological medicine (discussed in Chapter 1) might also explain the relinquishing of the perceived need to base diseases in specific pathology. A century ago, when the foremost medical sciences were all laboratory sciences, research primarily investigated the physiological or microbial basis of disease. A century later, epidemiology provides a way of investigating diseases, their causes, their outcomes, and their natural history that can remain agnostic to the underlying pathology, so long as diseases can be defined according to their measurable clinical manifestations. Epidemiology has an important—and perhaps unexpected—influence on the way we understand contemporary chronic diseases.

## 4.6. Conclusion

I defended a dispositional theory in which chronic diseases and mental disorders are generally bodily dispositions toward clinical findings. Dispositions are properties that manifest certain effects but that can remain unmanifested. This explains medicine's silent disorder phenomenon: disease dispositions are defined and diagnosed through their manifestations, yet can be asymptomatic or silent. For instance, mental disorders are characterized by diagnostic criteria that include symptoms and behaviors, but mental disorders can be in remission. The dispositional theory is better supported than alternatives advanced by philosophers, including alternatives that view diseases or mental disorders as processes, syndromes, or fictions. It also helps explain medicine's difficulty in explaining, predicting, and intervening in chronic diseases and mental disorders because different cases of the same kind of disposition can depend on different kinds of bodily pathology. On the other hand, the chronicity of chronic diseases and mental disorders partly reflects their incurability, which is a statement more about medicine than metaphysics (though cure may be more elusive when a disposition lacks a specific pathology). This chapter corroborates a finding repeated throughout Part 2 of this book, that scientific and metaphysical theories can aid our understanding of diseases of the new modern medicine.

# PART 3
# EVIDENCE

# 5

# Medical Evidence and Prediction


## Summary

It is now a truism that good medicine must be "evidence-based." But what is medical evidence? And how does it support medical predictions in prognosis and therapy? I first examine evidence-based medicine's measurement theory of medical evidence and its standard model of prediction, before defending an argument theory of medical evidence. The argument theory can help overcome several problems faced by the measurement theory, including problems with its approach to evidence evaluation and medical prediction.

Explanations in contemporary medicine are often theoretical or mechanistic, harnessing knowledge of underlying mechanisms or causes of disease. In contrast, predictions in the new modern medicine conform to a different model, an epidemiological model. I call this standard prediction model the Risk Generalization-Particularization Model. The first stage or inference in the Risk GP Model is Generalization, the extrapolation of an aggregate outcome from an epidemiological study population to a target population. The outcome is usually either an aggregate risk quantifying how common a clinical outcome is within a population, or it is a risk reduction or increase. The second inference is Particularization, or transformation of aggregate risks into patient probabilities. Both inferences rely upon substantive yet unarticulated assumptions. The Risk GP Model lines up squarely behind one of the two dominant prediction traditions in the history of medicine: the empiric tradition of predicting directly from experience with previous patients. Today the empiric tradition is represented by the evidence-based medicine (EBM) movement, which seeks to apply evidence from clinical epidemiology to patient care and places clinical epidemiology evidence atop an evidence hierarchy. Implicit in the hierarchy and the EBM literature is a measurement theory that views medical evidence as measurements of clinical outcomes, especially risks. Unfortunately, the measurement theory makes little room for diverse evidence that plays other roles in medical prediction, including providing support for the Risk GP Model's assumptions.



*The New Modern Medicine*. Jonathan Fuller, Oxford University Press. © Oxford University Press 2025.
DOI: 10.1093/9780190066178.003.0006

Thus I advocate for an inference theory or argument theory of medical evidence that understands evidence more broadly as facts or data that support premises in arguments for medical hypotheses.

## 5.1. From Explanations for Diseases to Evidence for Predictions

Part 2 of this book was about disease, especially theories of disease that help explain the occurrence of diseases or their special features. These include theories about the pathogenesis of contagious diseases and cancers (Chapter 2), a theory of the multifactorial etiology of contagious and noncommunicable diseases (Chapter 3), and a theory of the dispositional nature of chronic diseases and mental disorders (Chapter 4). Part 3 is about evidence in the new modern medicine, especially therapeutic evidence and therapeutic prediction; it is about the inferences used to predict how medical interventions will affect patients and the evidence used to support those predictions.

Physicians and scientists often explain the occurrence of a disease or its manifestations with reference to theories, models, mechanisms, or causes of health and disease. For instance, they explain how mutation or altered expression of cancer-related genes disrupts molecular mechanisms regulating somatic cell proliferation, how environmental etiologic factors interact with genetic predisposition to cause diseases like cancer or heart disease, and how interventions can work to prevent or treat these diseases by interacting with their etiology or pathology. Doctors could make forward-looking predictions in an analogous way, using causal models or models of mechanisms. However, there is a curious asymmetry in contemporary scientific medicine. Physicians use different kinds of reasoning to explain versus to predict. They prefer to use knowledge of causes and mechanisms to explain, and epidemiological evidence to predict. Of course, epidemiological evidence can also help establish causes and mechanisms of disease. But while these causes and mechanisms generally mediate medical explanations, they are often absent from medical predictions, which are often made more directly from empirical evidence provided by epidemiological research.

There are long medical traditions of making predictions from our understanding of mechanisms versus from direct empirical evidence. Nineteenth-century empirics bracketed theory and sought to establish prognosis and

therapeutics through direct clinical observation in the hospital, while nineteenth-century mechanists sought to exploit bacteriological causes and physiological mechanisms of disease learned through laboratory experiment. The dominant approach to medical prediction since the 1990s is that of EBM and is part of the empirical tradition. EBM was a movement in academic medicine that swept through the medical world, influencing medical journals, clinical guidelines, systematic reviews, evidence summaries, research reporting, and medical training. The fact that the adjective "evidence-based" is now so commonplace and EBM principles are so widely accepted speaks to the massive influence of EBM and reveals that EBM is no longer a movement pushing against medical orthodoxy but has become orthodoxy.

EBM changed how medicine evaluates evidence. One of the founding principles of EBM is that there is a hierarchy of evidence, and one should make diagnostic, prognostic, and therapeutic decisions using the best evidence according to the hierarchy (EBM Working Group 1992; Djulbegovic and Guyatt 2017). Epidemiological evidence is ranked higher in the hierarchy, while evidence from the laboratory or basic sciences is ranked lower. Lowly evidence of mechanisms is often said to support "pathophysiologic rationale" (EBM Working Group 1992) or "mechanistic reasoning" (Howick 2011), in which one reasons from hypothesized mechanisms of disease or medical intervention to make predictions about the effects of interventions.

In this chapter, I explore the more highly trusted predictions made from epidemiological evidence in EBM. These predictions standardly involve extrapolating from the results of a clinical epidemiology study such as a randomized controlled trial (RCT), an inference I examine further in Chapter 7. The results are standardly used to make inferences about population and patient risks, which I analyze further in Chapter 6. Risks quantify the frequency or probability of a clinical outcome for a population or for a patient. An RCT or a nonrandomized intervention study often estimates the decrease or increase in risk brought about by an intervention. One of the reasons that evidence from RCTs is so prized in EBM is that it is thought to be less prone to bias compared with observational epidemiological evidence. I explore the concept, implications, and evidence of biased evidence in Chapter 8.

Much of the philosophical literature on EBM has examined the virtues and vices of randomization and randomized trials (e.g., Worrall 2002, 2007; Cartwright 2007, 2010; Fuller 2019). Much of it has also critically discussed EBM's hierarchies of evidence. While my focus here will be the use of evidence

in medical prediction, the discussion does have relevance for these classic debates. For example, I introduce two problems with evidence hierarchies. First, the strength of epidemiologic evidence, including "gold standard" randomized trials, may be overstated. Second, EBM provides no approach for synthesizing or amalgamating diverse evidence. These problems are related to the way that evidence is understood and used in EBM.

EBM's framework for evidence and prediction is an epidemiological one, based in clinical epidemiology. In the second half of the twentieth century, clinicians created clinical epidemiology as a science that applies epidemiologic methods to clinical problems. Concerned that the results of epidemiology studies were often ignored by other clinicians, they then started the EBM movement. One of the movement's founders called EBM "the practical application of clinical epidemiology in patient care" (Daly 2005, 89). EBM is a major pillar of contemporary epidemiological medicine, in which the results of epidemiological research are integrated into clinical practice (Chapter 1).

Clinical epidemiology is a science of clinical measurement and quantification (Daly 2005), where the objective is often to estimate risks and treatment effects in a way that minimizes or at least accounts for the possibility of biased estimates. It is thus no surprise that EBM, as applied clinical epidemiology, inherits a measurement theory of medical evidence. In the measurement theory, evidence consists of measurements of clinical outcomes, especially risks, and predictions are made using these measurements. This theory explains why the EBM hierarchies value studies that minimize the risk of bias, but also why they struggle to accommodate evidence that doesn't directly estimate risks. In contrast, philosophers have articulated an argument theory for evidence (Cartwright and Hardie 2012; Fuller et al. 2024), in which medical evidence is that which supports premises in arguments about medical hypotheses, including medical predictions.

An argument theory of medical evidence can help overcome several problems faced by the measurement theory, including problems with its approach to evidence evaluation and medical prediction. Some of these problems will be discussed here; others will be discussed in later chapters of Part 3. In section 5.2, I locate EBM's origins in an empiric tradition of medical prediction. In section 5.3, I reveal its measurement theory of medical evidence. In section 5.4, I introduce its standard model of medical prediction, the Risk Generalization-Particularization Model. In sections 5.5 and 5.6, I discuss the assumptions made by Generalization and Particularization,

respectively. Then in section 5.7, I introduce the argument theory of medical evidence and argue that it can help overcome some weaknesses of the measurement theory.

## 5.2. Two Traditions for Medical Prediction

Physicians are predictors. Medical history has had its empirics, who predict directly from experience with previous patients, and its mechanists, who predict from theory. In this section, I illustrate these two traditions through historical examples and briefly trace the origins of EBM to the empiric tradition. Physicians today do not always recognize that they are fundamentally engaged in prediction. They use the term "prognostication" to describe the act of predicting future health outcomes but have no analogous term to describe the process of predicting the results of medical intervention. Thus I will first lay out some prediction concepts for medicine.

Scientific explanation has long been a fascination for philosophers of science, and rightly so. Much of research across the sciences can be understood in terms of the quest to explain what is going on in the world. However, medicine is a discipline especially concerned not only with explaining what is going on with the patient (why their toes are tingling, why they had a heart attack), but also with predicting what will happen to the patient. It is good to know why you had a heart attack, but better if you can predict whether you will have another heart attack in the future, and better still if you can predict what will prevent it from happening. In the philosophy of science, explanation has historically received more attention than prediction (Broadbent 2013).[1]

Broadbent (2013) identifies a couple of ambiguities in using the word "prediction" in science. He distinguishes two senses of the term: prediction as a claim, and prediction as an activity. Prediction activities are inferences; prediction claims are their conclusion. Medical forecasts, much like meteorological ones, are often probabilistic rather than definite. So I will count as prediction claims conclusions like the following: the probability that Priya will have a heart attack in the next ten years is 10%. Moreover, "prediction"

[1] Two early monographs in contemporary philosophy of medicine were titled *Discovery and Explanation in Biology and Medicine* (Schaffner 1993) and *How Scientists Explain Disease* (Thagard 1999).

can have a wide scope, including inferences about past, current, or future outcomes, or a narrow scope, including only future outcomes. In diagnosis, in which inferences are made about current outcomes, physicians speak about "clinical prediction rules" and the "positive predictive value" of a diagnostic test. However, because I am more concerned with prognosis and intervention in this chapter, I will focus on predictions in the narrower, future-facing sense.

We can call any representation of a recurring prediction activity or inference a "prediction scheme" or "prediction model." The classic point of departure for an explanation/prediction model in philosophy of science is the Deductive-Nomological (DN) Model of Carl Hempel and Paul Oppenheim (1948). The DN Model explains or predicts some physical fact (e.g., the acceleration of an object) by deducing it from initial conditions (e.g., the mass of the object and the force applied to it) along with laws of nature (e.g., Newton's Second Law):

Initial conditions
Laws of nature
Explanandum or prediction claim

All that distinguishes explanation from prediction in this model is that in explaining, the conclusion is a given (we start from below the line and work up), while in predicting, the premises are a given (we start from above the line and work down). The DN Model represents explanation and prediction as symmetrical activities, as being inferentially equivalent.

In a similar style, Paul Thagard (1999) argues that medicine is unified by a hierarchy of explanation schemes. Thagard represents medicine's most general explanation scheme as follows (modified slightly for presentation):

The patient is or has been subject to causal factors
The causal factors produce a disease with associated symptoms
The patient has the disease and symptoms

The patient's disease and symptoms are explained by citing causal factors that regularly produce diseases of that kind. Whether medical explanations typically conform to this general model is not my main concern. At least sometimes a physician will explain why a patient had a heart attack and its attendant symptoms (for instance) by citing causal factors like hypertension

that contributed (or might have contributed) to their heart attack. In the spirit of the DN Model's symmetry, contemporary physicians could in principle use this same causal reasoning to make a prediction. They could predict that a patient will have a certain disease from the known presence of certain causes. But without the support of epidemiological evidence, such reasoning would fail to live up to the ideals of contemporary medical prediction. Contrariwise, medicine's standard model of prediction is not used—and cannot be used—to explain.

Contemporary medicine's typical methods for explaining and its typical methods for predicting map onto different intellectual traditions in the history of medicine, which I will call the "mechanist" and "empiric" traditions.[2] Reasoning from causes or biomedical mechanisms of health and disease is more aligned with the mechanist tradition, which champions the establishment of medical theory or mechanisms of clinical causation. Meanwhile, contemporary prediction models are more aligned with the empiric tradition of predicting directly from past experience, which emphasizes observation, classification, quantification, and comparison.

Both traditions were alive in Hippocratic medicine, which influenced medical thinking right up to the nineteenth century. Humoral theory provided a general mechanistic model for predicting the results of interventions. Because the balance or harmonious mixture of the body's four humors constituted health while their imbalance or faulty mixture constituted disease, the doctor and patient could predict that an intervention which reverses the imbalance would cure disease and restore health. For instance, bloodletting could cure the excess of blood that was associated with fever. Meanwhile, Hippocratic prognostication was based on a simple empiric model of prediction: enumerative induction from past experience. The Hippocratics realized that particular diseases have a predictable natural course that can be learned from observing many cases. For instance, if previous patients had a characteristic fever and then either recovered or died after a fortnight, they could predict that the next patient encountered with the characteristic fever would recover or die after a similar number of days (Nutton 2012).

In empiric models of prediction, we predict directly from experience with previous patients. Notice that the Hippocratic prognostic model cannot explain why the patient will recover or die, which precludes any

[2] Drawing the same distinction, Bluhm and Borgerson (2011) and Howick (2016) call them the "rationalist" and "empiricist" traditions, respectively.

symmetry between prediction and explanation. It does not make sense to explain that the next patient will die because previous patients died. Their previous deaths are not the reason why the next patient will die. The explanation makes no more sense if the previous patients were followed within an epidemiological study and their outcomes quantified more precisely.

Although coexisting in Hippocratic medicine, the mechanist and empiric approaches were often warring historical traditions. Followers of Hippocrates founded the Dogmatic or Logical school of Ancient Greek medicine, which emphasized reasoning about the underlying causes of disease within the body and extended humoral theory with expanded anatomic knowledge. Later the Greek Empiric school favored empirical observation over theorizing about hidden causes. The Empirics based therapeutic predictions on trials with previous patients, relying on what appeared to work in the past, however strange or unpalatable the remedy was. Experience was often provided by case histories compiled by other physicians. In deciding whether to extrapolate from previous cases, they judged how similar those previous cases were to the new presenting case (Nutton 2012).

In the nineteenth century, when modern scientific medicine was emerging, the empiric[3] tradition was represented by "the numerical method" of Pierre Louis and his disciples. Louis famously applied the method to study the case records of seventy-seven patients with "pneumonitis" at La Charité hospital in Paris, comparing patients who were treated with bloodletting earlier in the course of disease with a matched group of patients who were treated with bloodletting later in the disease. Summarizing this application of the numerical method, Louis writes, "What was to be done in order to know whether bloodletting had any favorable influence on pneumonitis, and the extent of that influence? Evidently to ascertain whether, other things being equal, the patients who were bled on the first, second, third or fourth day of their illness, recovered more readily than those bled at a later period" (1836, 55). Louis's group comparison, with its notion of "other things being equal," was a precursor to modern epidemiological thinking. In fact, Alfredo Morabia

[3] Lester King (1991) surveys the many meanings of the term "empiricism" in nineteenth-century medicine, some of them associated with the numerical method or with the positive philosophy of Auguste Comte (2012). EBM's emphasis on measurement (discussed in the next section) has strong resonances with both the numerical method and Comte's positivism. However, EBM supporters deny that their movement is committed to any one philosophical theory of scientific evidence (including positivism) and argue that it instead draws on all major theories (Djulbegovic et al. 2009).

(2006) argues that Louis's method of reproducible clinical observation, medical quantification, and population thinking identifies him as the intellectual father of clinical epidemiology.

Meanwhile, Louis's French contemporary Claude Bernard helped to carry the mechanist flame. He advocated strongly for experimental laboratory science (especially physiology) as a basis for medical knowledge, claiming that "[e]mpiricism precedes science" (1875, 702–703). "What really should be done, instead of gathering facts empirically," writes Bernard, "is to study them more accurately, each in its special determinism. We must study cases of death with great care to try to discover in them the cause of mortal accidents, so as to master the cause and avoid the accidents" (Bernard 1865, 137). For Bernard, disease was a quantitative deviation from normal physiological function. Thus, by studying physiology and its disturbance in disease, the experimentalist could discover how to restore normal physiology and cure the disease.

As we saw in Chapter 1, Bernard's vision for a laboratory-based physiological medicine was well represented in the ideals for modern American medical education promoted in the early 1900s (Flexner 1910). On the other hand, William Osler, a transitionary figure in the making of modern medicine, was not only a great pathologist and champion of laboratory science; he was also known for promoting careful clinical observation. In their chapter on syphilis in the textbook *Modern Medicine*, Osler and Churchman endorse as a treatment regimen the method of "chronic intermittent treatment" because they believe it is "[t]he most rational plan, as well as the one which has given the best clinical results" (1907, 519–520). By "most rational," they mean the regimen most consistent with physiological and microbiological theory. By "best clinical results," they mean the regimen best supported by accumulated clinical experience. Mechanist and empiric justifications are invoked in the same sentence; there is no hierarchy of methods invoked justifying one over the other.

Although the laboratory sciences featured heavily in scientific medicine of this period, Louis's numerical method was not dead. The maturation of epidemiology and statistics as disciplines in the twentieth century allowed for more rigorous methods of population analysis and group comparison. But it was not until EBM placed epidemiological reasoning at the top of a "hierarchy of evidence" at the end of the twentieth century that the empiric approach to prediction officially surpassed the mechanist approach in priority. Miriam Solomon (2015) aptly characterizes EBM as "empiric medicine," a

distant relative of the Ancient Greek Empiric School. (Jeremy Howick 2016 draws a similar historical line.)

During the twentieth century, the science of epidemiology matured. However, it was a science of public health rather than of clinical medicine, investigating disease distributions and disease trends in human populations, as well as making inferences about the causes of disease from population data. Likewise, the influence of statistics on medical practice during this period was indirect. While statisticians led the statistically minded biometrical school in the early 1900s, their research penetrated biology more than medical thinking. The statistical foundations of the randomized experiment were developed by R. A. Fisher in the 1930s not in the context of clinical research but in agricultural research (Matthews 1995).

The 1900s were also a century of therapeutic reform, which provided the opening for the empiric tradition to seed medicine with epidemiology and statistics. Early in the century, therapeutic reformers of the American Medical Association sought to separate effective from quack remedies and turned to a combination of laboratory (animal, chemical) experiments as well as clinical experiments done in hospitals. Their clinical experiments were sometimes uncontrolled and sometimes included a control group matched to the treatment group based on clinical features judged to be relevant (Marks 1997). In 1948, a prototype for the modern randomized clinical trial was established by the British Medical Research Council's study of streptomycin for tuberculosis. (The U.S. Public Health Service ran its own, less famous randomized clinical trial.)[4] In 1970, the U.S. Food and Drug Administration began requiring randomized trials of new drugs, analyzed statistically, as a prerequisite for drug approval, and by the late 1970s most Western industrialized countries followed suit (Matthews 1995).

In a simple RCT, participants are randomly allocated to one of two or more study groups, and each group receives a different intervention, which may include a new or experimental intervention, an intervention currently in use, or a placebo or sham intervention (one meant to resemble the comparator intervention). According to Iain Chalmers (2005), epidemiologist Austin Bradford Hill's main justification for the use of randomization in the 1948 streptomycin trial (on which Hill was consulted) was to prevent foreknowledge among the study recruiters regarding the study group to which the next

[4] A 1931 trial of gold therapy may have been the first to use randomized treatment group allocation alongside double blinding of patients and investigators to group assignment (Lilienfeld 1982).

participant would be allocated. Otherwise, their foreknowledge might influence whether they admitted the next participant to the trial, which could taint the study results with bias (namely, "selection bias") by systematically influencing the composition of the study groups. In an RCT, participants, clinicians, and investigators may also be blinded or masked to the intervention that a participant is receiving.

According to Harry Marks, "By all apparent indicators, the second half of the twentieth century represents the 'statistical era' of clinical medicine" (1997, 129), especially with the rise of the RCT. Similarly, J. Rosser Matthews argues that after the RCT became the regulatory standard, "[t]he medical profession would now have to use statistics (a tool of the social sciences) just as much as it had formerly used the laboratory (a tool of the natural sciences). This transformation marked a major shift differentiating nineteenth from twentieth century medical thinking in the realm of medical therapeutics" (1995, 148). However, this "transformation" occurred in therapeutic evaluation and regulation rather than immediately in medical thinking. Ultimately, it was clinical epidemiology and EBM that brought these developments into the clinic.

Robert Fletcher, coauthor with Suzanne Fletcher of the first clinical epidemiology textbook, describes the setting for a new discipline's appearance: "We grew up in a very biomedical era. . . . [E]veryone was a laboratory scientist; and it seemed to me that the kind of science that they brought to bear on patient care, which was mainly logical argument from laboratory data on mechanisms of disease, just wasn't the best possible way of answering these clinical questions" (quoted in Daly 2005, 21).[5] Similarly, Alvan Feinstein objected to the scarcity of clinical research involving live patients compared to basic laboratory science at medical conferences. Feinstein envisioned a science of medicine, an alternative to the laboratory, in which clinical practice would be studied empirically and subjected to quantitative

[5] I owe subsequent details of the history of clinical epidemiology and EBM to Jeanne Daly's (2005) research, which includes interviews with many of the founders of both disciplines. While it is told here as a Canadian and American story, there were parallel developments in the U.K. (Daly 2005). The Scottish physician Archie Cochrane championed the need for RCTs, trial registries, and trial summaries, particularly with his influential 1972 book *Effectiveness and Efficiency*. The British obstetrician Iain Chalmers argued in favor of systematic reviews of the literature for entire medical specialties, starting with his own. In 1993, Chalmers helped found the influential Cochrane Collaboration (now known as Cochrane), an international network of reviewers that produces systematic reviews for all areas of healthcare. Finally, in addition to the *Journal of the American Medical Association*, the *British Medical Journal* was one of the mainstream medical journals that embraced EBM early on.

analysis. He called this science "clinical epidemiology." Just as epidemiology studied public health, clinical epidemiology would study clinical practice. In 1967, David Sackett became chair of the first department of clinical epidemiology and biostatistics, which was established at McMaster University in Canada as a department in the school of medicine. Sackett saw clinical epidemiology as the application of epidemiology to diagnostic, prognostic, and therapeutic questions. The clinical epidemiologist's faith in clinical observation and quantification as the foundation for a science of medicine echoed the appeals of medical quantifiers in the nineteenth century and therapeutic reformers earlier in the twentieth.

The Framingham Heart Study, a cohort study begun in 1947, was probably the first large-scale "follow-up" or prognostic study in clinical epidemiology (Giroux 2011). The Framingham study was not only public health research; it was also clinical research (Giroux 2012): the investigators conducted clinical interviews, clinical examinations, and diagnostic testing, including bloodwork and electrocardiograms, and the research identified clinical risk factors that could be measured in the clinic and generated algorithms for the quantification of a patient's risk of heart disease—all with the aim of guiding patient care.

Meanwhile, the Cholesterol Treatment Trialists meta-analysis of clinical trials (Baigent et al. 2005) is a paradigm example of therapeutic research in clinical epidemiology. Meta-analysis is a method of pooling the results of multiple statistical studies that answer the same question, a method advanced by Thomas Chalmers (among others) in the U.S. in the 1980s. It is a favorite tool of clinical epidemiology, used to synthesize the results of clinical trials and inform medical treatment, often as part of a systematic review of the medical literature. The Cholesterol Treatment Trialists meta-analysis pooled data on thousands of patients enrolled in randomized trials from the 1990s and 2000s to determine the safety and effectiveness of cholesterol-lowering statin drugs meant to prevent heart attacks and strokes.

What distinguished the Framingham Study and the Cholesterol Treatment Trialists meta-analysis from other epidemiological or statistical studies is that they were seen as answering clinical questions of prognosis and treatment, respectively. But medicine still had to listen to the answers. A concern among clinical epidemiologists at McMaster that medicine was largely oblivious to the clinical research evidence motivated the EBM movement in the early 1990s. EBM urged the use of clinical epidemiology evidence in medical practice. Guyatt, who coined the term "evidence-based medicine," described

it as a "new paradigm," "the practical application of clinical epidemiology in patient care." According to Guyatt, "One needs rules of evidence that are, essentially, clinical epidemiology. . . . [R]easoning on the basis of physiology often proves misleading without empirical testing" (quoted in Daly 2005, 89).

The mission statement of EBM has received many formulations, but probably the most quoted is provided by Sackett and colleagues: "Evidence-based medicine is the conscientious, explicit and judicious use of current best evidence in making decisions about the care of individual patients. The practice of evidence-based medicine means integrating individual clinical expertise with the best available external clinical evidence from systematic research" (1996, 71). It is difficult to object to the style of practice represented by this definition, but the details matter. For instance, it does not reveal what kind of external clinical evidence EBM considers to be the "best evidence." For that, we must look to the EBM hierarchies of evidence. Jeremy Howick (2011) as well Robyn Bluhm and Kirstin Borgerson (2011) argue that the hierarchy of evidence is the central tenet of EBM. It posits that RCTs (the "gold standard" of evidence) and especially systematic reviews of RCTs are better grounds for determining treatment effectiveness than observational epidemiologic studies, and that both kinds of evidence are superior to evidence from either basic biomedical research or clinical experience. Hierarchies of evidence were initially promoted as a tool to help clinicians select evidence for decision-making but soon became used primarily as evidence grading schemes for medical guideline development. EBM also developed standards for research reporting (including clinical trial reporting) and systematic review.

On close inspection, EBM is essentially the application of the statistical tools, concepts, and theories of the clinical epidemiology tradition—plus the hierarchy of evidence and a few other "rules of evidence"—to clinical reasoning and clinical care. Its founders called EBM a "paradigm shift" (EBM Working Group 1992), and perhaps that was prophetic. EBM made a name for itself early by publishing "User's Guides to the Medical Literature" in the *Journal of the American Medical Association* as well as journals and other resources devoted to evidence summaries. In 1994, Sackett founded the world's first EBM unit, the Oxford Centre for Evidence-Based Medicine (Daly 2005). In less than ten years, EBM was quickly and enthusiastically embraced by the major medical journals, professional societies, Cochrane and other groups for evidence evaluation and synthesis, and clinical guideline development

groups; it was also incorporated into medical training as a "new approach to teaching the practice of medicine" (EBM Working Group 1992).

Much of the discussion of EBM in the philosophy of medicine is—not surprisingly—about evidence in EBM. However, different kinds of evidence lend themselves to different models of prediction. Evidence of biomedical mechanisms is linked to models of mechanistic prediction, a point that Howick (2011) seems to acknowledge by including "mechanistic reasoning" (in which one reasons through a mechanism for a clinical outcome) in his version of the hierarchy of evidence. Similarly (I argue), epidemiological evidence is used within an *epidemiological* prediction scheme, a new kind of empiric model. With the elevation of clinical epidemiology evidence (especially evidence from randomized trials) to "gold standard" status, the epidemiological model has quietly become the standard model in contemporary healthcare. Yet it receives no formal treatment, little explicit discussion, and no name. I call it the Risk Generalization-Particularization Model, and I introduce it in section 5.4. But first, I will examine EBM's pyramid of evidence and the theory of medical evidence concealed within its walls.

## 5.3. Evidence in Evidence-Based Medicine

EBM arrived at the end of the 1900s with a zeal for setting medicine on a foundation or "evidence base" formed from clinical epidemiology research. In the new modern medicine, clinical epidemiology evidence became synonymous with evidence. The most iconic image of EBM is its hierarchy of evidence for treatment effectiveness, which represents clinical epidemiology evidence as the apex of an evidence pyramid or ladder. The RCT, the "gold standard," is the star that shines at the top. EBM's privileging of epidemiological research has attracted much criticism, some of which I will rehearse. Not discussed, however, is that the hierarchy embodies a theory of evidence, one that sees medical evidence as measurements of clinical outcomes, especially risks.

Hierarchies of evidence appear throughout EBM guides, clinical guidelines, and journal publishing policies. However, the origin of the hierarchy has not been rigorously studied. Jeanne Daly (2005) reports a story repeated by EBM founders, in which Sackett invented the hierarchy of evidence to resolve an intractable dispute at a medical consensus conference. The earliest example of a hierarchy of which I am aware is in a report of the

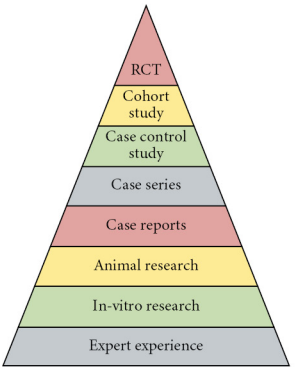
Canadian Task Force on the Periodic Health Examination (CTFPHE) from 1979, of which Sackett was an author. Now known as the Canadian Task Force on Preventive Heath Care, the Task Force was a pioneer in rating evidence and developing evidence-based clinical guidelines. In 1979, it made recommendations on the use of diagnostic and therapeutic interventions during periodic health exams (health checkups). It provided a system for grading the quality of evidence backing an assessment of the effectiveness of an intervention, from I (highest quality) to III (lowest quality):

I One or more RCTs
II-1 Cohort or case-control studies
II-2 Comparisons across time or location, or uncontrolled experiments with dramatic results
III Expert opinion based on clinical experience, case studies, or expert committee reports. (CTFPHE 1979)

The first mention of the expression "gold standard" to describe the RCT appeared a few years later, in 1982 (D. Jones and Podolsky 2015).

Since then, there has been a burst of evidence hierarchies put forward by different groups, often differing in terms of what kinds of evidence they name and include. There are also different evidence hierarchies for diagnosis, prognosis, and treatment. Figure 5.1 contains two illustrative treatment hierarchies.

(A)




(B)

| Quality of evidence | Study design | Lower quality if | Higher quality if |
|---|---|---|---|
| High | Randomised trial | Study limitations<br>–1 serious<br>–2 very serious | Large effect<br>+1 large<br>+2 very large |
| Moderate | | Inconsistency<br>–1 serious<br>–2 very serious | Dose response<br>+1 evidence of a gradient |
| Low | Observational study | Indirectness<br>–1 serious<br>–2 very serious | All plausible confounders<br>+ Would reduce a demonstrated effect or<br>+ Would suggest a spurious effect when results show no effect |
| Very low | | Imprecision<br>–1 serious<br>–2 very serious | |
| | | Publication bias<br>–1 likely<br>–2 very likely | |

**Figure 5.1** Evolution of the hierarchy of evidence for treatment from earlier forms (A) to later forms (B).
Djulbegovic and Guyatt (2017, 417).

Although they may appear different, and often do contain important differences, every hierarchy is similar in listing various evidence types. Evidence types are distinguished according to their source, which is typically a kind of research study, whether experimental (RCT) or non-experimental ("observational study"), human-based (epidemiological) or laboratory-based (animal research, in vitro research). A hierarchy assigns evidence to a "level" according to its "quality" or "certainty." In EBM, higher quality evidence generally supports stronger conclusions about which interventions are effective as well as stronger recommendations about their use. Higher quality evidence is generally less likely to be biased, where research bias is now understood as deviation from the truth brought about by systematic errors arising from study design. (I discuss the concept of bias further in Chapter 8.) For instance, randomization (when properly concealed) is thought to rule out selection bias, a family of biases resulting in systematic differences among study groups due to how the groups are formed (Howick 2011). Randomization can block selection bias by allowing study group formation to be decided by the least systematic of all selectors: random chance.

Some hierarchies make finer-grained distinctions among study types than others. The hierarchy in Figure 5.1A distinguishes cohort studies, case-control studies, case series, and case reports, while the hierarchy in 1B lumps these under the heading "observational study." Some hierarchies include additional types of evidence, such as systematic reviews or meta-analyses of RCTs, which are placed above a single RCT. Finally, some hierarchies, like the one in Figure 5.1A, are "categorical" (J. Stegenga 2014): their rankings hold for all evidence tokens of a given type (e.g., RCTs are always higher quality than cohort studies). Other hierarchies, like the one in Figure 5.1B, which represents the popular GRADE approach (Balshem et al. 2011), are non-categorical and allow upgrading and downgrading of quality assessments so that observational studies may sometimes rise above RCTs in the ranking.

However, in all hierarchies, non-epidemiological evidence, including evidence from the laboratory, either sits below epidemiological evidence (Figure 5.1A) or is not overtly included (Figure 5.1B), reflecting EBM's empiric stance toward prediction. Moreover, RCTs always sit above observational studies in hierarchies of evidence for treatment effectiveness (provisionally in non-categorical hierarchies). Finally, all hierarchies identify the best evidence (a single study or body of similar studies) so that the evidence user can use it to make predictions. Hierarchies do not encode a system for integrating or synthesizing different kinds of evidence. According to EBM

gurus Benjamin Djulbegovic and Gordon Guyatt (2017), the hierarchies of evidence reflect "the first EBM epistemological principle." The basis for this principle is that "not all evidence is created equal, and that the practice of medicine should be based on the best available evidence" (416).

This privileging of certain kinds of evidence (epidemiological studies, RCTs) over others and the exclusive use of only "the best available evidence" have been a common source of complaints in philosophy about evidence hierarchies.[6] The first problem with evidence hierarchies is an overconfidence in evidence at the top of the hierarchy, which, it is argued, is not as reliable as it may seem. Commentators who point to the privileging of RCTs often quote from the textbook *Evidence-Based Medicine*: "If the study wasn't randomized, we'd suggest that you stop reading it and go on to the next article in your search. . . . Only if you can't find any randomized trials should you go back to it" (Strauss et al. 2011, 70). While non-categorical hierarchies qualify this position, it can be argued that too much faith is still bestowed on RCTs. Many worries have been raised within medicine and philosophy about exaggerated views of the reliability of RCTs.[7] I will not repeat all these criticisms here, but the problem of unreliable and biased evidence will reappear in Chapter 8. Some of the criticisms are methodological, questioning the value of randomization, placebo controls, frequentist statistical tests, and pooling of trial data, for example. Others are social, concerning, for instance, corrupting incentives in many trials, including industry sponsorship. Ironically, many early uses of the "gold standard" label in the medical literature were in critiques of the epistemic hegemony granted to RCTs (D. Jones and Podolsky 2015). In sections 5.5 and 5.6, I add to these concerns by showing that *even when* RCTs are unbiased, they may not yield reliable predictions.

The second problem with evidence hierarchies is that they recommend to evidence users what they consider to be the best available evidence, to the exclusion of other evidence. Many philosophers and other critics have argued that we should instead be considering multiple kinds and lines of evidence in concert.[8] Different kinds of evidence have different strengths and

[6] For a defense of the hierarchy of evidence, see Howick (2011). For a review of criticisms, see Bluhm and Borgerson (2011).

[7] For examples from philosophers, see Urbach (1985), Worrall (2002, 2007), Borgerson (2009), La Caze (2009), Solomon (2015), J. Stegenga (2018), and Deaton and Cartwright (2018). For selected responses on the virtues of randomization, see Howick (2011), Fuller (2019), Larroulet Philippi (2022), and Martinez and Teira (2024).

[8] For examples from philosophers, see Bluhm (2005), La Caze (2011), Andersen (2012), Clarke et al. (2014), Parkkinen et al. (2018), Anjum et al. (2020a), J. Stegenga (2018, 2022), and A. Park et al.

weaknesses and provide complementary information. Multiple lines of evidence fit together like pieces of a puzzle or are mutually reinforcing. Another EBM guidebook, *Users' Guides to the Medical Literature*, suggests that "the totality of evidence" should be examined. But it goes on to clarify that "[i]n practice, this means that our inferences (and decisions) are best informed by systematic reviews (ie, syntheses of the totality of relevant *high-quality evidence*)" (Guyatt et al. 2015, 16, my emphasis) and that EBM provides principles to help identify high-quality evidence. Systematic reviews often include only RCTs, even when nonrandomized evidence is available, and guidelines typically grade their recommendations solely based on the highest-level evidence available. Evidence from higher in the hierarchy is typically not integrated with evidence lower down, and especially not with non-epidemiological evidence.[9] I will argue that EBM's theory of evidence contributes to this problem.

While philosophers and physicians have picked at some of the epistemic assumptions inherent in the hierarchy, especially the assumptions underlying the ranking of different kinds of evidence, another assumption has gone virtually unnoticed: What is evidence in EBM? What makes the results generated by the studies listed in the hierarchy "evidence"? Ironically, for a movement that calls itself "evidence-based medicine," EBM has said very little about the sense in which epidemiological studies provide evidence for medicine. The *Users' Guides to the Medical Literature* offers what its authors consider to be "a broad definition of evidence: any empirical observation or report of a symptom or mental state constitutes potential evidence" (Guyatt et al. 2015, 16). Unfortunately, this definition doesn't tell us when potential evidence becomes evidence, or what distinguishes evidentiary from nonevidentiary empirical observations.

(2023). For selected responses on the problems with evidence of mechanisms, see Howick (2011) and Howick et al. (2013).

[9] The *Users' Guides* suggests that theory has a role in EBM (Guyatt et al. 2015), but this role is not well developed. Similarly, while knowledge of mechanisms is mentioned in the GRADE literature, GRADE provides sparse advice on how to evaluate or integrate it (Perillat and Mercuri 2022) and recommends against upgrading or downgrading quality of evidence ratings directly using mechanistic rationale (Guyatt, Oxman, Kunz, et al. 2011). The integration of effect estimates from randomized and nonrandomized intervention studies is a developing area of guidance for GRADE (Cuello-Garcia et al. 2022). One exception to EBM's silence on integrating dissimilar evidence is its long-standing recommendation to integrate relative treatment effects with untreated risks to calculate absolute treatment effects (Strauss et al. 2011).

While EBM has not been very explicit about the nature of evidence, implicit in the hierarchy is the theory that medical evidence consists in measurements of clinical outcomes, especially risks (including both untreated risks and risk reductions or increases). The term "risk" has different meanings, even in epidemiology and medicine, where it can refer to the hazards or harms posed by an exposure or intervention. In Part 3 of this book, I use "risk" to refer to concepts having more to do with the measurable frequency or probability of medical outcomes. These are usually harmful outcomes but can also be desirable outcomes.

After evaluating the quality of the evidence and before "applying" the evidence to the patient, the next step in the EBM process is to ask, "What are the results?" (Guyatt et al. 2015). Results typically take the form of an aggregate clinical outcome in a population. In prognostic research such as a cohort study, it is common to represent the outcome, be it the occurrence of disease or death, using the absolute risk (AR). The AR in epidemiology is the proportion of patients in the population with the clinical event (AR = number of events/number of people). In therapeutic studies such as RCTs, the outcome in question is the effect size of an intervention, how much of an effect the intervention has in the population. The effect size is commonly represented using outcome measures derived from the AR that measure a risk reduction or increase. One such outcome measure is the absolute risk reduction (ARR), the difference between the AR in each study group ($ARR = AR_{control} - AR_{treatment}$). Another outcome measure is the relative risk reduction (RRR), which is the ratio of the ARR to the control group's absolute risk ($RRR = ARR/AR_{control}$).

In EBM, the results of an epidemiologic study do not primarily support a belief, the rejection or acceptance of a hypothesis, or the confirmation or disconfirmation of a scientific theory. The primary function of studies is to measure clinical outcomes, especially risks, and epidemiologic studies are seen as providing the best estimates or measurements of these outcomes.[10] The results of these studies, the "evidence," are measurements. For instance, therapeutic evidence is a measurement of an intervention's effect, often expressed quantitatively as a risk reduction or a risk increase. Epidemiologists Kenneth Rothman and Sander Greenland summarize this view of epidemiological evidence: "[S]cientific evidence can usually be

[10] Ross Upshur et al. (2001) and Leah McClimans (2013) comment on the shift toward clinical measurement with EBM.

viewed as a form of measurement. If an epidemiologic study sets out to assess the relation between exposure to tobacco smoke and lung cancer risk, the results can and should be framed as a measure of causal effect" (2005, S150). They note that like any measurement, this one can suffer from measurement error. Here the errors include those relating to study design, uncontrolled confounding, and bias.

To measure an outcome in a population is to assign a value of a metric (or "measure") to that outcome (the "measurand"). According to Norman Bradburn and colleagues (2017), only when the measure is properly coordinated with the concept used to refer to the measurand as well as our procedures for applying the measure can medicine achieve good measurement results. In EBM, the measurand is a health or disease outcome represented by concepts like death (in prognosis) or the effect size (in therapy). Commonly used population-level measures include the AR and the ARR. Finally, the measuring device or procedure is typically an epidemiologic study like a cohort study or RCT. According to the hierarchy, some of these devices provide more reliable or accurate measurements than others, owing to a lower threat of study bias.

It is thus no surprise that EBM makes ample use of measurement constructs. For example, in stating the "first EBM epistemological principle," Djulbegovic and Guyatt write, "[T]he higher the quality of evidence, the closer to the truth are *estimates* of diagnostic test properties, prognosis, and the effects of health interventions" (2017, 416, my emphasis).[11] Estimation involves assigning values of a measure to a measurand, such as assigning an ARR of 3% to the effect size. Again, the role of research studies in EBM is to estimate or measure outcomes, and evidence consists in the results of these measurements.

Where are these outcomes being measured? The most direct measurement is made in the sample study population, the individuals who are observed and counted. However, the point of making measurements in a study population is to provide a basis for inferences about some other—usually wider—population. EBM thus sees evidence as providing measurements

[11] Similarly, on the standard approach to predicting cardiovascular disease using prognostic calculators, clinical guideline authors write, "By its nature, such an approach requires a platform for reliable quantitative *estimation* of absolute risk based on data from representative population samples" (Goff et al. 2014, S53, my emphasis). In the approach to rating therapeutic evidence developed by the GRADE group, the aim is to grade "*estimates* of the effects of health care interventions" (Guyatt et al., 2015, 11, my emphasis).

of patients that the clinician encounters in practice, albeit measurements that are less direct. They are less direct in that they are mediated by further assumptions, which we will explore in sections 5.5 and 5.6. Even the more direct measurement of outcomes in the sample population requires certain assumptions. For instance, measuring the effect size in an RCT involves making a causal inference, which relies on both empirical and metaphysical assumptions. (See Fuller 2019 for a discussion of various assumptions.) Most (if not all) scientific measurements are mediated by theoretical assumptions. Measuring heart rate by palpating the radial pulse assumes a connection between ventricular contraction and arterial pulsation; measuring body temperature with an old-fashioned mercury thermometer assumes a predictable relationship between body temperature and the expansion of mercury.

In evaluating evidence, the hierarchies assist us in judging the "quality of evidence" or "certainty of evidence" (Guyatt et al. 2015; Hultcrantz et al. 2017). This concept refers to the certainty attached to the measurement of the effect provided by evidence rather than (for instance) the extent to which the evidence supports a hypothesis about the patient's response to treatment. Further measurement constructs such as "valid" and "precise" are used to describe good evidence (Guyatt et al. 2015). Validity here refers to measurement accuracy, while precision refers to measurement error. A valid, precise study result is one in which the measurement estimate is unbiased or close to the truth, and it has a narrow confidence interval. The confidence interval reflects the random measurement error one can expect. Or, as the textbook *Evidence-Based Medicine* puts it, "The confidence interval (CI) provides the range of values that are likely to include the true risk and quantifies the uncertainty in *measurement*" (Straus et al. 2011, 85, my emphasis).

The very idea that we can compare the quality of different kinds of evidence side by side and select the best evidence for answering a clinical question assumes that all the relevant evidence is performing a similar function (that of estimating the measurand). Only then does it make sense to say that an RCT does a better job than a nonrandomized study or a laboratory study in answering the question *What is the clinical effect?* When various strands of evidence are viewed as providing redundant information, and one of these strands is stronger, it makes some sense to rely on this strongest strand, especially in cases where the other evidence disagrees in its estimation of the effect. Moreover, it makes some sense to exclude studies that do not directly provide this information. In answering the question *What is the clinical effect?*, it makes sense to exclude studies such as laboratory research that do not

directly measure an intervention's clinical effects by comparing differently treated clinical populations.

EBM's measurement theory is perhaps unsurprising, given the origins of EBM in clinical epidemiology, a discipline concerned from the beginning with measuring and quantifying clinical outcomes (Daly 2005). Of course, measurements are made throughout the sciences, often in the process of generating supporting evidence. A chemist might measure various properties of a chemical sample to provide evidence for its purity. In the clinic, a physician might measure all sorts of parameters on physical exam and bloodwork to provide evidence for a diagnosis. In public health epidemiology, an epidemiologist might measure the correlation between an environmental exposure and cancer (even quantifying the relationship using population risks) in providing evidence that the exposure causes cancer. In all these cases, the *measurement serves as evidence* for a hypothesis.

In contrast, most of the work of evidence evaluation in EBM is assessing measurement accuracy and error, which gives the impression that *evidence serves as a measurement* rather than the other way around. While these measurements subsequently function as inputs in predictions, as I will suggest in the next section, these predictive inferences are often implicit and far less attention has been given to them in EBM compared to measurement, especially within evidence hierarchies. Granted, in evidence hierarchies the "quality" of evidence is synonymous with our level of "confidence" or "certainty," an epistemic concept often used to describe degrees of belief or support for hypotheses or theories. Yet the confidence or certainty to which the hierarchies refer is confidence in our measurements.[12]

Spencer Hey (2015) argues that clinical trials have multitudinous theoretical implications. (Angus Deaton and Nancy Cartwright 2018 agree.) Not only do trials confirm or disconfirm a "theory" about the effectiveness of the specific treatment for the specific disease in the specific study population. They also provide evidence for effectiveness with respect to wider classes of therapies, patients, and diseases; they provide evidence for theories of how the treatment works; and they provide evidence for theories about disease pathophysiology. I agree that medical evidence has much wider and more

[12] Guyatt et al. write, "The GRADE approach involves rating our confidence in *estimates* of the effects of health care interventions (also referred to as quality of evidence) as high, moderate, low, or very low" (2015, 11, my emphasis). The GRADE Working Group later clarified that this confidence refers to our certainty that the effect size exceeds a certain threshold value (the threshold for recommending the treatment) or falls within a certain range (Hultcrantz et al. 2017).

diverse import than is typically recognized in EBM. My concern is that these kinds of evidential implications are being missed in EBM, and they are being missed because EBM implicitly views medical evidence as measuring outcomes in populations of patients. It should not surprise us that EBM is relatively unconcerned with the wider theoretical implications of evidence because EBM favors empiric prediction and regards reasoning about mechanisms and theory as speculative.

EBM's measurement theory contributes to the first problem with evidence hierarchies that I discussed earlier: the reliability of epidemiologic evidence may be exaggerated, especially compared to non-epidemiologic evidence. When good evidence is seen as that which reliably measures clinical outcomes, especially risks, non-epidemiologic studies are at an immediate disadvantage. Risks and other clinical outcomes are the kinds of outcomes that epidemiologic studies measure by design. A clinical trial measures the effect size using the RRR or ARR by literally comparing two populations of patients, tallying clinical events, and calculating a ratio or difference. In comparison, laboratory studies do not literally compare populations of people, and so are ill suited to measuring the RRR or ARR of clinical events. They are better suited to building models of biological mechanisms. Mechanistic information can provide important evidence for medical prediction (as we will see). But this role does not consist in measuring clinical outcomes, and thus this role goes unrecognized by the evidence hierarchies.

The issue is that EBM's theory of medical evidence seems to be derived from the kinds of results that epidemiologic studies provide. A considered theory of medical evidence should instead be guiding the interpretation of these studies as a source of evidence for medicine. EBM was founded by clinical epidemiologists who were impressed by the ability of clinical epidemiology studies—especially RCTs—to rule out sources of error to which other kinds of evidence may be more susceptible. For instance, laboratory studies may be less reliable in making clinical predictions because they do not investigate live patients. The "first EBM epistemological principle," the hierarchy of evidence, thus gives much more space to epidemiologic evidence. (In fact, non-epidemiologic evidence is not even listed in the GRADE hierarchy in Figure 5.1B and is infrequently discussed in the GRADE literature.) Moreover, owing to the disciplinary bias of EBM, evidence appraisal has narrowly focused on the kinds of results that epidemiologic studies provide: aggregate clinical outcomes. EBM's theory of medical evidence emerges from this framework.

The other problem with evidence hierarchies introduced earlier is that hierarchies instruct their users to use "the best available evidence" and they encode no approach for synthesizing qualitatively different kinds of evidence. Many evidence hierarchies list systematic reviews, which may include a meta-analysis, a tool for combining evidence. However, meta-analysis is designed to pool similar statistical results only; for example, a meta-analysis of RCTs pools only RCT results. A meta-analysis will often not combine RCT results with the results of nonrandomized studies. And because non-epidemiologic studies typically do not measure aggregate clinical outcomes, they cannot be included in a meta-analysis of RCTs that is pooling aggregate clinical outcomes. Therefore, EBM's theory of evidence as consisting in risk measurements deters systematic reviewers from combining epidemiologic evidence with non-epidemiologic evidence, throwing up a conceptual barrier to their synthesis.

Of course, if measuring risks is a worthy role for research studies, then the fact that non-epidemiologic studies fail to perform that role is a problem for them, not for the hierarchy of evidence. But the risks in question are, in the first place, aggregate outcomes in study populations, and medical evidence should inform clinicians in the care of individual patients outside of those studies. It must be shown how these measurements are relevant for patient care. If these measurements are not directly relevant, then it is a mistake to exclude non-epidemiologic studies merely for failing to provide them, especially if these studies have a role to play in establishing the relevance of epidemiologic measurements.

In summary, EBM views medical evidence as measurements of clinical outcomes, especially untreated risks as well as risk reductions and increases. The job of making these measurements relevant to patient care is done by EBM's standard model of prediction, to which I now turn.

## 5.4. The Risk GP Model

The rise of EBM was a victory for the empiric tradition. Empiric methods of prediction are now accorded a special status in healthcare. Because the evidence in EBM is epidemiological, it necessitated a new epidemiological model of prediction: the Risk Generalization-Particularization Model. The Risk GP Model is the standard model in evidence-based prognosis and intervention, but it receives little discussion as a prediction scheme in healthcare.

Here I will reconstruct it, before examining its assumptions in the subsequent two sections.

We can think of the Risk GP Model as having two versions, one for prognosis and one for intervention. I will start by describing the prognostic version. Simple induction from past experience like that used in Hippocratic prognostication is sometimes represented in the hierarchy of evidence, but typically in the basement level. EBM cautions doctors against relying on this evidence: their recollections are likely to be unsystematic and to suffer from cognitive biases. They should instead rely on epidemiological research that systematically quantifies experience with previous patients.

Enter the absolute risk. Two examples of the use of the absolute risk in prognosis are the American College of Cardiology/American Heart Association (ACC/AHA) cardiovascular risk calculators (Goff et al. 2014; Khan et al. 2023) and the Framingham Risk Score (Anderson et al. 2016). These algorithms calculate the ten-year absolute risk of cardiovascular events—heart attacks or strokes—in individuals without established disease, using similar risk factor variables as inputs.[13] For forty-five-year-old men, the Framingham calculator tells us that the "normal" risk (in patients without too many risk factors) is 5.6%, meaning that on average, 56 out of 1,000 of them will develop a cardiovascular event over ten years. As we will see in a moment, an important threshold for treatment recommendations is a ten-year risk of 10% (Mangione et al. 2022). The physician can use these algorithms to calculate the absolute risk of cardiovascular events for groups of patients in their practice.

When they do so, they tacitly rely on the Risk GP Model. The absolute risk computed by these algorithms was validated in the finite cohorts of patients in which they were developed (Khan et al. 2023). The physician's practice population will not consist in any of these cohorts of people. (Several of the cohorts were followed decades ago.) Thus the physician must extrapolate the absolute risk—measured in the study cohort—to their target practice population. In keeping with the language used by EBM, I call this prediction activity Generalization.[14]

[13] These variables include sex, age, systolic blood pressure, smoking, diabetes mellitus, total cholesterol, and HDL cholesterol.

[14] The term "generalization" unfortunately makes it sound like the predictor is extrapolating to the "general population," to all patients. The target population could be the general population but need not be.

Even once they have "generalized" the absolute risk measurement to their target population, the physician faces an added puzzle. The absolute risk literally describes the proportion of outcomes in a population, an aggregate outcome. The direct question the physician wants answered is whether the patient sitting on their examination table will have a cardiovascular event over the next ten years. To say that 10% of people of the patient's age, sex, and cardiovascular profile will have an event is not a direct answer to that question. How, then, does the physician descend from the population level to the level of the individual patient?

The answer lies in a distinct concept of medical risk. Doctors and guidelines speak of "the patient's risk" of the outcome, which is often taken to be equal in value to the absolute risk in a target population in which the patient can be located. Because the absolute risk is about the population and not the individual, we have good reason to suspect that conceptually speaking the patient's risk is related to the absolute risk but is distinct. Namely, the patient's risk is the probability that the patient will have the outcome. (I will have more to say about the meaning of this probability in Chapter 6.) This probability is inferred from the absolute risk. Making use of the tight link between frequencies and probabilities, the physician (implicitly) transforms the absolute risk into their patient's probability.[15] If the absolute risk is 10% in the population, the probability that the patient will have a cardiovascular event is 10%. To contrast this second inference with Generalization, I call it Particularization.

To some extent, the ACC/AHA acknowledges that risk prediction relies on these inferences: "By its nature, such an approach requires a platform for reliable quantitative estimation of absolute risk based on data from representative population samples. It is important to note that risk estimation is based on group averages, which are then applied to individual patients in practice" (Goff et al. 2014, S53).

[15] It is more challenging to interpret the relevance of an aggregate measure for an individual when the patient-level outcome is nonbinary, such as quality of life or severity of disease. However, researchers often get around this problem by dichotomizing the patient-level variable so that the absolute risk can be computed. For instance, the primary outcome in a recent meta-analysis of antidepressants was a reduction in total score on a depression rating scale of at least 50%, while a secondary outcome was full-stop symptom remission (Cipriani et al. 2018). It is then possible to calculate the absolute risk of a 50% score reduction as well as the absolute risk of remission (though this study calculated the odds rather than the absolute risk).

When we combine Generalization and Particularization in sequence, we end up with the Risk Generalization-Particularization (Risk GP) Model. We can represent the Risk GP Model for prognostic predictions as below.

Risk GP Model for Prognosis
In the study population, the absolute risk of outcome *O* is *r*
...
In target population *F*, the absolute risk of outcome *O* is *r*
...
For patient *a* in target population *F*, the probability of *O* is *r*

The Risk GP scheme is a serial inference in which the conclusion of the Generalization (first inference) is the first premise in the Particularization (second inference). Both inferences rely upon suppressed premises or assumptions that are represented by the ellipses. I will scrutinize these assumptions in the next two sections.

First, we must consider how the Risk GP Model makes predictions about medical intervention. To predict how their patient will respond to medical intervention, a clinician could call upon their own experience with similar patients. The difficulties here are many and include the challenges of making causal inferences from data that are unsystematic, incomplete, and at high risk of bias (or so the EBM critic warns). Thus their EBM guides suggest that they turn to an epidemiological study—preferably an RCT or systematic review of RCTs—that measures the effect for them.

In the 1980s, the U.S. Preventive Services Task Force (USPSTF) adopted the evidence rating methodology developed by the Canadian Task Force (1979). USPSTF is a leader in the development of evidence-based guidelines on medical prevention. For instance, it periodically publishes guidelines on the use of statins, cholesterol-lowering drugs that have been widely prescribed for the prevention of cardiovascular disease since the 1990s (Greene 2007). The 2022 USPSTF guidelines on the use of statins in the primary prevention of cardiovascular disease (prevention in people without a history of cardiovascular disease) recommend prescribing a statin for people ages forty to seventy-five years with one or more cardiovascular risk factors (dyslipidemia, diabetes, hypertension, or smoking) and who have an estimated cardiovascular disease risk of 10% or more. They further recommend selectively prescribing a statin for similar people who have a cardiovascular disease risk of 7.5% to 10% (Mangione et al. 2022). USPSTF

suggests using an established ACC/AHA risk calculator (Goff et al. 2014) to sort patients into cardiovascular risk categories. (A new ACC/AHA risk calculator has since been published; Khan et al. 2023.)

USPSTF gives these two recommendations a B grade and a C grade, respectively. The grades are based on USPSTF's having moderate certainty in the evidence that statins will have an overall net benefit in cardiovascular disease prevention. (The B grade is due to statins having a "moderate net benefit" for those with 10% or greater risk; the C grade is due to statins having a "small net benefit" for those with a 7.5% to 10% risk.) The evidence comes from a systematic review and meta-analysis of the epidemiological literature on the benefits and harms of statins conducted by USPSTF (Chou et al. 2022). The review included twenty-three RCTs measuring benefits and harms as well as three observational studies measuring harms. It largely found no increased risk of harm with statins. Meanwhile, the review found a decreased risk of all-cause mortality, stroke, heart attack, and cardiovascular events overall. The guideline predicts that by prescribing a statin, a physician can achieve a relative risk reduction in cardiovascular events of about 30%.[16] If the absolute risk is 10% among untreated patients, it is then 7% among statin-treated patients.

What does this prediction mean? It means that as an effect of the statin, the physician can predict that for every hundred patients treated, three fewer will have a cardiovascular event over the next ten years on average compared to their untreated patients. Once more, the physician has made an inferential leap. The RRR of 30% was measured in the fixed population included in the meta-analysis, but the physician is predicting the RRR for their patients. Whether or not the physician realizes it, their prediction requires Generalization of the effect measured in the study population to their target population of interest.

Of course, the physician (hopefully) does not line up their patients, give all of them a prescription for a statin, and tell them that 30% of those who were destined to have a cardiovascular event (whoever those may be) will now have their event prevented by the drug. Instead, they will often tell the

[16] All-cause mortality: RRR = 8% (95% confidence interval [CI]: 13% to 2%); fatal or nonfatal stroke: RRR = 12% (95% CI: 32% to 10%); fatal or nonfatal heart attack: RRR = 33% (95% CI: 40% to 25%); composite cardiovascular events: RRR = 28% (95% CI: 36% to 19%) (Chou et al. 2022). The composite cardiovascular event endpoint varies among RCTs included in the meta-analysis and often differs compared to the composite cardiovascular event outcome calculated by the ACC/AHA risk calculator. For a skeptical and personal take on the results of statin trials from a philosopher of science, see Worrall (2021).

individual patient sitting on their examination table that a statin will "lower their risk" by 30%. They are again referring to the patient's individual risk or probability of the outcome, which they hope will be reduced in some sense by the medication. The inference from the effect size in the target population to the change in the probability of the outcome for the patient is another Particularization. If the relative risk reduction is 30%, then the probability that the patient will have a cardiovascular event is 30% lower in relative terms given statin treatment compared to no treatment (e.g., 7% rather than 10%).

Although they do not use the terms Generalization or Particularization, EBM experts Paul Glasziou and David Mant explain that "[a]pplying results to treatment decisions" involves two stages: "The first stage involves an assessment of the transferability of the trial evidence [to the target population]; the second deals with the application by the clinician to the individual" (2007, 88–89).[17]

The Risk GP Model for therapeutic predictions is represented below.

Risk GP Model for Intervention
In the study population, the effect size for outcome *O* due to intervention *X* is *r*

...

In target population *F*, the effect size for outcome *O* due to intervention *X* is *r*

...

For patient *a* in population *F*, the change in probability of *O* due to *X* is *r*

The physician extrapolates the effect size equal to *r* measured in the study population to the target population they have chosen (first inference), and then derives a change in probability of the outcome equal to *r* for their patient in the target population (second inference). Both inferences require assumptions. Most of the debate around clinical trials in EBM has not centered on these inferences and their assumptions but rather on whether clinical trials provide the best evidence. However, the evidence that they provide is usually only direct evidence for the first premise in this chain of reasoning. It is not only in prevention and primary care that risk calculators and predictions from clinical trials are standard. An example from oncology is the PREDICT model for predicting breast cancer mortality risk (Candido

[17] David Kent and colleagues write, "The use of evidence from clinical trials to support decisions for individual patients is a form of 'reference class forecasting': implicit predictions for an individual are made on the basis of outcomes in a reference class of 'similar' patients treated with alternative therapies" (2018, 1).

dos Reis et al. 2017) and the Early Breast Cancer Trialists' Collaborative Group's (2021) meta-analysis of RCTs of trastuzumab for the treatment of early stage, HER2-positive breast cancer. This meta-analysis measured the reduction in risk of breast cancer recurrence and death from breast cancer from the biologic trastuzumab. Peter Keating and Alberto Cambrosio (2011) argue that oncology as a medical specialty developed around the cancer clinical trial after the 1960s.

The Risk GP Model does not discriminate among disease cure, disease prevention, and disease management. A change in risk could refer to a disease that the patient already has, so that the treatment reduces the risk of still having the disease in the future (cure). The change in risk could instead refer to a disease that the patient does not yet have, in which case the treatment reduces their risk of developing the disease (prevention). Or for patients with the disease, it could refer to some manifestation the treatment is designed to prevent or some beneficial outcome the treatment is designed to bring about (disease management). In all cases, "the patient's risk" is the target of treatment in some nebulous sense that cries out for analysis (which will have to wait until Chapter 6, in which I examine both the concept of a preventive and curative medical intervention and the concept of the patient's risk).

I capitalize "Generalization" and "Particularization" to signify that these terms refer to the very specific inferences we have been discussing. In a broader sense, generalization could mean extrapolating any information from a study, while particularization could mean individualizing this information in some way. You will not find the Risk GP Model explicitly laid out in textbooks of medicine or clinical epidemiology. Yet in applying risk measurements from studies to patients, healthcare providers implicitly call upon it.

The Risk GP Model is the standard model of prediction in prognosis and treatment in that it is held up as a normative ideal by EBM authorities, if only implicitly. In other words, it is the *gold standard*. The *Users' Guides to the Medical Literature* claims that applying study results to patient care requires asking whether you can "generalize" or "particularize" the results to your patient (Guyatt et al. 2015, 5).[18] Furthermore, Risk GP is the prediction scheme

[18] Other textbooks, surveys, and articles also identify generalization and individual risk assessment as the gold standard approach (Szklo and Nieto 2007; Post et al. 2013; Fuller 2013b; Glasziou and Irwig 1995; S. Goodman 1999; Djulbegovic et al. 2011; Straus et al. 2011).

upon which evidence-based practice guidelines rely, as we have just seen.[19] Practice guidelines set the *standard of care* in medicine, another sense in which this scheme is the standard model of medical prediction. I have been using cardiovascular risk prediction and statin-mediated risk reduction to illustrate the model. But in prognosis, Risk GP can accommodate any study that measures the absolute risk of any outcome, and in intervention, Risk GP can handle any therapeutic study that uses a measure of effect size derived from the absolute risk.

The Risk GP Model is EBM's standard model, the empiric tradition's epidemiological prediction scheme. It involves two inferences: extrapolation (Generalization) of epidemiological study results to a target population, and transformation (Particularization) of the results into probabilities about an individual patient. Its inferences often go unrecognized and their assumptions unnoticed. In the next two sections, I will more closely examine each of its inferences, Generalization and Particularization.

## 5.5. Generalization

In 2013, the ACC/AHA rejected the widely used Framingham Risk Score for predicting cardiovascular disease risk partly because it was derived in an exclusively White cohort from a restricted geographic area. ACC/AHA developed its own calculator, the Pooled Cohort Equations, which was validated in more diverse cohorts, to expand the "generalizability" of its risk estimates (Goff et al. 2014, S71). Eventually, the generalizability of this calculator was also questioned, with validation studies showing overprediction of cardiovascular events in some populations and underprediction in others, partly due to changing population patterns of risk factors and treatment (Khan et al. 2023). So, ten years later, ACC/AHA released a new risk calculator, arguing that because of the many cohorts in which it was derived and validated, its estimates "are broadly generalizable to the target population of interest" (Khan et al. 2023, 1990).

Turning to treatment effect predictions, the USPSTF statin guidelines (Mangione et al. 2022) are careful to match the target population to the trial

[19] I surveyed clinical guidelines recommending the use of common prescription medications and found that generalization from clinical trial results was used every time to support treatment recommendations for a guideline's target patient population (Fuller 2013b).

population and are more cautious than the ACC/AHA guidelines in one respect. Unlike ACC/AHA (Grundy et al. 2019), USPSTF recommends statins in primary prevention only for people with certain risk factors (diabetes, smoking, dyslipidemia, or hypertension) because patients in the statin trials had these risk factors. USPSTF also concludes that the evidence is insufficient to recommend a statin for those over seventy-five years of age because older patients are underrepresented in statin trials. In this section, I will introduce the problem of extrapolating from epidemiological studies, which is examined in more detail in Chapter 7.

First, it is worthwhile mentioning a prior problem that confronts extrapolations from intervention studies. Measuring the effect size in an intervention study is not as easy as counting outcomes. It requires a causal inference, in which the investigator or evidence-user concludes that the association found in the study—quantified using the ARR, RRR, or some other measure—is an effect of the treatment rather than an association with some other explanation. This causal inference is given much thought in philosophy of science, epidemiology, and EBM. (For overviews, see Steel 2011; Deaton and Cartwright 2018.) What matters for the present discussion is that a causal inference is required prior to our causal prediction, yet our causal prediction is a separate inference that makes its own assumptions.

In clinical epidemiology and EBM, "internal validity" describes the accuracy of measurements made of the study population. In contrast, "external validity" refers to the extent to which the study results accurately predict the results in some target setting. Evidence hierarchies sometimes consider external validity in evaluating the quality of evidence. For instance, the GRADE hierarchy includes "indirectness" of the effect measurement as one consideration that should lead an evidence appraiser to downgrade their confidence in the effect estimate (Guyatt, Oxman, Kunz, et al. 2011). A study is "indirect" when, for instance, the study population is unrepresentative of the target population.[20] This suggests that the effect estimate in question is an estimate or measurement of the effect in the target population made by collecting data from a study population. If the study population is a sample from the target population, then it makes some conceptual sense to consider the study results as measurements or estimates of values in the target.

[20] It can also be indirect when there are important differences in the intervention or outcome, or when the study does not compare the two interventions under consideration head-to-head (Guyatt, Oxman, Kunz, et al. 2011).

However, not all study populations are subsets of a larger target population. Even when a study population is a subset, GRADE evaluates indirectness by deciding whether dissimilarities between study and target are biologically relevant. So the inference from study to target is better conceived as an extrapolation from one population to another rather than a statistical inference from a sample. It is conceptually awkward at best to frame the results in the study as a measurement of the results in the target and indirectness as a determinant of measurement accuracy. In section 5.7, an argument theory of evidence reconceives of study results as supporting a premise in an inference about the target.

Further limiting the GRADE approach to external validity, GRADE uses the method of simple extrapolation to assess indirectness for the study population, suggesting that an evidence user should rate down the quality of the evidence only when there is an uncommon and compelling biological reason to think that any differences between the study and target populations will result in a substantial difference in the magnitude of the relative treatment effect (e.g., the relative risk reduction). I argue against simple extrapolation in Chapter 7. GRADE's overall approach is to consider whether any study-target differences are important and likely to make a difference to the outcome. While GRADE provides some useful suggestions and examples, it offers no general solution to the problem of when differences make a difference.

Several authors in the medical literature have suggested strategies for extrapolating from epidemiological studies. (For a review, see Post et al. 2013.) These strategies typically consist in a series of questions that one may ask to probe whether a study result is "generalizable." For instance, in an article titled "How to Generalize Efficacy Results of Randomized Trials: Recommendations Based on a Systematic Review of Possible Approaches," Piet Post and colleagues list six questions, including "Are there patient differences that may diminish the treatment response?" and "Are there important differences in patient compliance that may diminish the treatment response?" (2013, 642). These questions provoke further ones: *What kinds of patient differences may diminish treatment responses?*, and *What differences in patient compliance are important?* Once again, in order to answer the questions that Post and colleagues pose, we need to determine what kinds of similarities and differences matter in general, and then develop more rigorous methods for investigating them that link these general considerations to domain-specific considerations.

The challenge of transporting a conclusion we made about the study population to some other target setting is often called "the problem of extrapolation" (Steel 2008). The problem of extrapolation asks *How can we determine whether the same result we found in the study will be found in our target?* One way to understand extrapolation is as an inference like Generalization. We can then examine the assumptions we need to make this inference valid. The problem of extrapolation becomes partly a problem of articulating and exploring these assumptions. Without doing so, Generalization will remain elliptical, just as we modeled it in the previous section.

For Generalization, the assumption in question is a "representativeness assumption" or "sufficient similarity assumption," stating that the study population and conditions are representative of (or sufficiently similar to) the target population. How representative or how similar we need the study to be depends on how accurate we expect our prediction to be. No matter how demanding we are, articulating the assumption requires understanding the kind of similarity that is needed for extrapolating the study result we have in mind. In general terms, the required similarity is causal. Because the sufficient similarity assumption relies on causal details, even empiric approaches to medical prediction cannot completely ignore causal-mechanistic considerations. I will articulate the sufficient similarity assumption in Chapter 7. For now, I will consider two examples that illustrate the kind of causal-mechanistic information the assumption invokes (but at a more general and abstract level).

Andrew Park and colleagues (2023) analyze the failed deployment of efavirenz in Zimbabwe as an example of how mechanistic evidence is relevant in extrapolating from RCTs and how the EBM hierarchies err by failing to integrate it with the results of RCTs. Efavirenz is an antiretroviral agent used to treat chronic HIV infection that was shown in RCTs done in the U.S. to safely and effectively lower HIV viral load. In 2015, the Zimbabwean government, following the guidance of the World Health Organization, rolled out efavirenz across the country. Unfortunately, many Zimbabweans who took the drug experienced significant neuropsychiatric side effects and stopped using it (L. Nordling 2017). The reason the safety and efficacy measured in the RCTs did not generalize to Zimbabwe was a widespread prevalence among Zimbabweans of a variant of a liver enzyme (CYP2B6) that metabolizes efavirenz. The variant is less efficient, leading to an accumulation of efavirenz in the brain (Nyakutira et al. 2008). The Zimbabwean population is not sufficiently similar to the American population due to this causal difference. This conclusion was not merely a post-hoc determination

made on postmortem, so to speak, but was predicted ahead of time by a Zimbabwean geneticist whose warnings—based on an understanding of efavirenz's mechanism of action—were apparently ignored (Park et al. 2023).

Another example is provided by a more recent emerging infectious disease: COVID-19. In late 2020 during the COVID-19 pandemic, several RCTs of vaccines against SARS-CoV-2, the virus that causes COVID-19, were completed. mRNA vaccine trials showed a relative reduction in symptomatic infection of approximately 95% (Baden et al. 2021; Polack et al. 2020). These trials had a three-month follow-up and were conducted at a time when the original or wild-type variant of SARS-CoV-2 was dominant. As the pandemic continued, there emerged many SARS-CoV-2 variants, mutant strains of SARS-CoV-2 that often had a slightly different structure. Immunological studies showed that SARS-CoV-2 neutralizing antibodies from vaccinated individuals bound to certain variants with similar efficiency as they bound to the wild-type variant, but that they bound to other variants (especially the Omicron variant) with reduced efficiency, suggesting the potential for "immune escape." Moreover, after five or six months, the levels of neutralizing antibodies in the blood of those vaccinated were reduced, especially among older people. Together, these findings showed a time- and variant-dependent waning of neutralizing antibody titers against the virus (Lassaunière et al. 2022; J. Evans et al. 2022).

Because antibody neutralization is a central component of the causal mechanism of vaccine-induced viral immunity and a correlate of immune protection, these kinds of laboratory findings led health experts to predict reduced effectiveness against infection among vaccinated populations in which individuals were more than five or six months out from completing vaccination and populations in which variants demonstrating some immune escape in vitro were dominant (F. Collins 2021; Tabak 2022). Both the initial neutralizing antibody levels generated after vaccination as well as the structure of the wild-type SARS-CoV-2 variant (specifically the structure of its "spike protein" that binds to human cells) are causal details that should be included in the sufficient similarity assumption for the COVID-19 vaccine trials. Because they differed along these two causal dimensions, the populations described above were potentially not sufficiently similar to the trial populations to extrapolate the effect size. Indeed, observational studies done in these populations suggested reduced vaccine effectiveness against infection (Tartof et al. 2021; Andrews et al. 2022).

Generalization requires a sufficient similarity assumption, that the study and target are similar in causally relevant ways. We are not done with the problem of

the external validity of trials; I will return to it in Chapter 7. For now, I will move on to the second of the Risk GP Model's inferences.

## 5.6. Particularization

How does the Risk GP Model come up with anything to say about individual patients when it extrapolates aggregate outcomes from population studies? The ACC/AHA is explicit that these aggregates are "applied to individual patients." They pause on the confusing meaning of aggregate risks for individuals: "This process is admittedly imperfect; no one has 10% or 20% of a heart attack during a 10-year period." They also note the related difficulty of understanding the relationship between aggregate outcomes and patient-level prevention: "Individuals with the same estimated risk will either have or not have the event of interest, and only those patients who are destined to have an event can have their event prevented by therapy." They offer no conceptual way out of these problems and simply conclude that no matter what approach the clinician takes, "the clinician must apply the average results obtained from groups of patients to the individual patient in practice" (Goff et al. 2014, S53). In Chapter 6, I examine the concepts of a preventive or curative medical intervention and of the patient's individual risk. Here I will examine this "application of the average results to the individual patient" and the assumptions it makes.

As epidemiologists Benjamin Djulbegovic, Iztok Hozo, and Sander Greenland observe, "The usual practice is simply to identify single-case probability with the frequency in the smallest identified class ('reference set') in which the patient can be placed" (2011, 310). They nicely summarize the Risk GP Model's solution to the epistemological problem of crossing the void between populations and individual patients: the physician infers patient probabilities from finite frequencies (absolute risks). Aggregate clinical outcomes are provided by Generalization and are relative to the target population ("reference set") to which the physician extrapolated. When the aggregate outcome is the absolute risk, the physician derives the patient's probability of the outcome. When the aggregate outcome is the relative or absolute risk reduction, they infer the change in the patient's probability of the outcome.[21] I called these inferences Particularization.[22]

[21] Absolute risks are relative to a target population. Accordingly, patient probabilities are *conditional* probabilities, conditional on the target population.

[22] Philip Dawid (2017) calls the inference from group frequencies to individual probabilities "individualization" or the "group to individual (G2i) inference," and the resulting individual risk an "individual*ized* risk."

The subtle transition from absolute risks to patient risks often goes unnoticed, and thus Generalization and Particularization are sometimes presented as a single step, as when the *Users' Guides to the Medical Literature* asks of the physician, "[C]an you generalize (or, to put it another way, particularize) the results to your patient?" (Guyatt et al. 2015, 5). The same medical sources that discuss Generalization typically do not describe the inference from aggregates to individuals.[23] However, Particularization is a distinct inference with two unique assumptions, both of which are seldom expressed.

The first indispensable assumption, the "population membership assumption," assumes that the patient under consideration is a member of the reference or target population. Otherwise, there are no grounds for measuring the probability of the outcome for the patient from the absolute risk of the outcome in that reference population. Careful work is needed to establish that this assumption is true, especially if the target population is defined according to a medical diagnosis, which is often the case. As the physician is typically less than 100% certain of a diagnosis, they should be less than 100% certain of this assumption—perhaps much less so. Their uncertainty as to the diagnosis and this assumption should influence their uncertainty as to whether the patient will have the outcome or whether the intervention will benefit the patient.[24] For example, their uncertainty as to whether the patient has diabetes should strongly influence their uncertainty as to whether the patient will have a heart attack and whether antidiabetic drugs or a statin will help them. Unfortunately, in medicine uncertainty is often equated with an epidemiological risk, which is inferred from group data and does not reflect how confident the physician is in whether the patient is a member of the group.

Particularization's second assumption, the "unbiased population assumption," is that the probability that a patient will have their prognosis assessed is independent of their risk of the outcome and the probability that they will respond to the intervention in a particular way. Individuals who would

[23] Yet many epidemiological and medical authorities use the probabilistic concept of "the patient's risk." Glasziou and Irwig write, "To identify patients who should expect benefit to be greater than harm, we need to predict each patient's risk" (1995, 1358). Writing on risk and probability, Stephen Goodman notes that we "take the property of a group—for example, the percentage of patients that dies—and call that our patient's risk. But when we transfer a number from a group to an individual patient, its meaning can qualitatively change" (1999, 604).

[24] The proof of this principle is provided by the total probability equation: $p(O) = p(O|F)p(F) + p(O|\neg F)p(\neg F)$. We can express the physician's uncertainty using credences, represented by the probabilities in this equation. A physician's uncertainty as to whether the patient will have the outcome ($p(O)$) depends not only on the patient's probability relative to the target population $F$ ($p(O|F)$), but also on the physician's uncertainty as to whether the patient is a member of $F$ ($p(F)$).

develop the clinical outcome or would benefit from the intervention are neither more likely nor less likely to have their treated and untreated prognosis assessed. Then the probability that a patient will have the outcome is equal to the proportion of individuals in the population who will have the outcome.[25]

We can model this situation with an urn containing light gray and dark gray balls (see Chapter 6), each ball numbered consecutively from 1 to 100. If you draw a ball blindly from the top of the pile, the probability that you have drawn a given ball (say, ball number 37) should be equal to the probability that you have drawn any other ball—but only if the draw is "fair." The draw will be fair if the "chance setup" is unbiased (Hacking 2001), for instance, if balls are randomly mixed and the mechanism for selecting them is not influenced by their color (as in a lottery). In this case, the probability of a light gray ball is equal to the proportion of light gray balls in the urn. On the other hand, if all the light gray balls are deliberately placed near the top of the pile and you are aware of it, the probability that you have drawn, say, ball number 37, where ball number 37 is known to be light gray, is *not* equal to the probability that you have drawn ball number 73, where ball number 73 is known to be dark gray. In that case, the probability of a light gray ball is not equal to the proportion of light gray balls.

To complete the analogy, if there is no reason to suspect that a patient presenting for cardiovascular risk assessment from among the target population is any more or less likely to have a heart attack, then the probability that they will have a heart attack can be supplied by the absolute risk. However, if (as seems plausible) patients with angina-like chest pain are more likely to have their cardiovascular risk assessed, and angina-like chest pain is known to be positively correlated with future heart attack, then the probability that patients with angina-like chest pain will have a heart attack should be higher than the absolute risk when the target population is defined using classic cardiovascular risk factors (which do not include angina-like chest pain).

In summary, Particularization relies on two assumptions. The population membership assumption holds that the patient under consideration is a member of the reference or target population. The unbiased population

[25] It is most reasonable to interpret these probabilities epistemically, as the clinician's credence. There may well exist an objective correlation between risk assessment and the outcome or between receiving intervention and intervention responsiveness. However, as long as the clinician does not know about the correlation, it is rational for them to assign an equal credence in the outcome to each patient in the target population.

assumption holds that the probability that a patient will have their prognosis assessed is independent of their risk of the outcome and the probability that they will respond to the intervention in a particular way.

The assumptions of Generalization and Particularization are unarticulated in medicine and will sometimes fail without our noticing. If any one of the assumptions fails, the physician's prediction may fly far of the mark. To mix metaphors, an inferential chain is broken if there is a gap in even one of its links. My concern here is relevant to one of the problems with evidence hierarchies that I introduced earlier, an overconfidence in study results that are ranked highly by the hierarchies. Even if an RCT measurement is unbiased, it may not yield reliable Risk GP predictions because these predictions rely on additional assumptions that are not supported by the RCT results but that require additional evidence. The decision to accord study results the status of "high-grade evidence" should consider the use to which that evidence will be put in addition to the accuracy of the study's measurement.

## 5.7. An Argument Theory of Medical Evidence

I argued that the EBM evidence hierarchies view evidence as consisting in measurements of clinical outcomes, especially risks. Risks then serve as inputs in EBM's standard model of prediction, the Risk GP Model. Prediction is the last step in the EBM workflow (Guyatt et al. 2015), often described as "applying evidence to the individual patient." For EBM, applying medical evidence often amounts to predicting patient risks from epidemiological measurements. In this section, I will provide an alternative to the measurement theory, an argument theory of medical evidence, and argue that it can help overcome several problems.

Prognostic and therapeutic inferences are predictive inferences. Physicians predict what will happen to a patient on versus off intervention. These inferences should, in some sense, be "evidence-based," and by comparing outcomes in groups of patients on versus off intervention, comparative epidemiologic studies, including RCTs, certainly provide relevant evidence. But this evidence, taking the form that it does of an estimate of the effect in a study, directly supports only the first premise in a chain of inference that I explored in the previous three sections. This inference involves further empirical premises that require further evidence. Predictive inferences should be evidence-based all the way through, not only at the front end.

The kinds of evidence that the Risk GP Model requires include support for a sufficient similarity assumption, which establishes causal similarity between the study and target population. They include support for a population membership assumption, which establishes that the patient is a member of the target population. And they include support for an unbiased population assumption, which establishes that the probability that a patient will have their prognosis assessed is independent of their risk of the outcome and the probability that they will respond to the intervention in a particular way. The evidence needed to support these assumptions might take many forms and may be experimental or non-experimental, clinically based or laboratory based. But it cannot be provided by simply running another clinical trial to produce another measurement of the effect. In other words, EBM's standard model of prediction requires evidence beyond measurements of risks in a study population.

This need for further kinds of evidence is a challenge for the evidence hierarchies and the theory of medical evidence that they suggest. When evidence is viewed as consisting in measurements of risks, and the quality of evidence is influenced mostly by the accuracy or unbiasedness of these measurements, then these other kinds of evidence and the ways in which their quality should be assessed are easily ignored. A theory of evidence as epidemiological measurement might be suitable for clinical epidemiology, if we conceive of clinical epidemiology as a science of clinical measurement, but it is inadequate for scientific medicine. EBM needs a more permissive theory of medical evidence, one that recognizes the diverse roles that medical evidence plays in prognostic and therapeutic inferences (as well as in diagnostic inferences).

A promising alternative is provided by Nancy Cartwright and Jeremy Hardie's (2012) "theory of evidence for use." Robyn Bluhm and Kirstin Borgerson (2011) observe that an explicit theory of evidence is lacking in EBM, while Rani Anjum and colleagues (2020a) call for expanding the notion of evidence in EBM. In *Evidence-Based Policy*, Cartwright and Hardie provide an expansive theory of evidence specifically for evidence-based policy, but their theory is easily extended to EBM. Part of their theory involves a general view of evidence and evidential warrant, and part of it elaborates causal considerations that are relevant to predicting the effectiveness of policies. The latter considerations are relevant to Generalization and will be explored in Chapter 7. For now, I will introduce their theory of evidence, apply it to medical predictions, and defend it against the measurement theory of medical evidence.

Cartwright and Hardie (2012, 19) summarize their theory of evidence as follows: "A well-established empirical claim e is evidence for h if and only if e can be rendered as a premise in an argument, A, and A is a good argument for h; or, e is a premise in A′, where A′ is a good argument for a premise in a good argument, A, for h; and so forth." Evidential relevance is a three-place relation involving e, A, and h: e becomes evidence for a hypothesis when it warrants that hypothesis via a specific logical argument with the hypothesis as its conclusion. Because most scientific arguments involve multiple distinct empirical premises, a hypothesis typically receives its warrant from a combination of multiple distinct lines of evidence.

According to their theory, evidence is a "well-established empirical claim" (Cartwright and Hardie 2012, 19). What distinguishes evidence from nonevidentiary reasons that one might also find in an argument (such as normative ethical claims) is that evidence is empirical; that is, "evidence claims report facts about the world" (21). The view that evidence is a premise in an argument might be closer to the philosopher's idea of evidence than the scientist's.[26] In the evidence hierarchies, evidence refers to the results of studies or to the clinician's experience rather than to propositions. We could modify Cartwright and Hardie's concept of evidence slightly so that evidence refers to the data or facts that are reported by evidence claims rather than the claims themselves. Or we could understand their theory as offering a concept of evidence claims rather than evidence. While I am inclined toward the former option and will write about evidence as consisting in facts or data, it matters little for their theory because the main contribution of their theory is that it helps us distinguish well-established facts (or claims) that are evidence from well-established facts (or claims) that are not. The results of an RCT might make an empirical claim well-established, but what makes them evidence for the hypothesis that a patient will benefit from the intervention studied is that they support a premise in an argument for that hypothesis.

By referring to evidence claims as "well-established," Cartwright and Hardie do not mean to restrict evidence to the results of RCTs; they suggest that evidence will sometimes render evidence claims merely probable, thus leaving room for evidence from all sorts of studies, not only those that are most likely to provide an unbiased measurement. Further, they allow that

[26] Among four concepts of evidence discussed by Peter Achinstein (2001), the idea that evidence is a proposition that justifies a hypothesis is most closely affiliated with what Achinstein calls epistemic situation evidence, or "ES-evidence."

evidence can be indirectly relevant to a hypothesis by supporting a premise in a sub-argument for a premise in the main argument for the hypothesis. While they focus on what we can call "confirming evidence," or evidence *for* a hypothesis, a friendly addition to their theory would be an account of disconfirming evidence, evidence *against* the hypothesis. On their theory, evidence for h (confirming evidence) supports a premise in an argument (or sub-argument) for h. By extension, evidence against h (disconfirming evidence) would either support a premise in an argument for a conclusion denying h, or it would undercut a premise in the argument for h.[27]

Cartwright and Hardie's theory of evidence is reductive in the following sense. Rather than espousing a unique logic of evidential warrant, a confirmation theory such as hypothetico-deductivism or Bayesianism, their theory reduces evidential warrant to logical implication, essentially leaving it to the sciences of logic and mathematics to fill in the details.[28] Consequently, their theory tells us more about what makes something evidence than about how much evidence it provides, or how good or strong the evidence is.[29] Moreover, their theory is pluralist in that it admits different kinds of logic (deductive and inductive) and argument schemes rather than committing to a universal logic or universal argument scheme for evidence-based reasoning.[30] The theory starts from local argument schemes made in particular scientific domains and then identifies the empirical premises in need of evidence.

[27] Julian Reiss's (2015) "pragmatist theory of evidence" similarly recognizes both "direct" and "indirect" evidence. However, Reiss understands indirect evidence differently, as that which is incompatible with an alternative hypothesis. Reiss contrasts his theory with an "experimentalist theory of evidence" in which hypotheses receive justification by being tested through a reliable gold standard study design or through a study that mimics the gold standard to some degree. He associates the experimentalist theory with EBM, and the pragmatist theory with the public health epidemiology evidence from decades ago showing that smoking causes lung cancer. Reiss (2014) analyzes several other theories of evidence from philosophy of science, which include Cartwright and Hardie's (2012), Achinstein's (2001), and several classic confirmation theories. Upshur and Colak (2003) discuss an approach to medical evidence grounded in argumentation and informal logic as a corrective to EBM's restrictive view of evidence.

[28] For example, according to one version of Bayesian confirmation theory (Howson and Urbach 2006), Bayesian conditionalization is the logic of induction, and e is evidence for/against hypothesis h if and only if $p(h|e) \neq p(h)$.

[29] Reiss (2014) calls these two features of evidence "support" (what makes something evidence) and "warrant" (how much confirmation a body of evidence provides). Reiss argues that an adequate theory of evidence must address both and thinks that Cartwright's fails to adequately address support because it leaves it to the sciences to tell us what makes for a good argument. I won't weigh in on whether we should expect the theory to go further, but I'll note that Cartwright and Hardie (2012) do go further by discussing arguments used in evidence-based policy and general principles for their evaluation.

[30] John Norton (2021) argues that there is no universal inductive logic.

Adapting their theory to medicine gives us an "argument theory" or "inference theory" of medical evidence. Rather than serving as a measurement of an aggregate clinical outcome, medical evidence as evidence supports a premise in an argument for a hypothesis (or, in the case of disconfirming evidence, undercuts the premise or supports a premise in an argument for a conclusion that denies the hypothesis). Evidence is *medical* evidence (as opposed to policy evidence, for instance) when it helps warrant a *medical* hypothesis such as a diagnostic, prognostic, or therapeutic claim. Examples of medical hypotheses include "Priya has diabetes," "Priya will have a cardiovascular event over the next ten years," and "A statin will prevent a cardiovascular event for Priya." Because these medical hypotheses are prediction claims, medical evidence generally supports premises in predictive inferences. These predictive inferences can be sub-arguments within practical or ethical argumentation supporting decisions, which creates an avenue for medical evidence to provide (partial) justification for medical decisions, policies, and recommendations via the argument theory.

On the argument theory, facts or data that support premises in arguments for medical hypotheses are medical evidence, regardless of the source of the facts or data and whether they are generated by an epidemiological study. In contrast, on the measurement theory, facts or data are medical evidence only if (and whenever) they are measurements of clinical outcomes, which generally excludes facts and data generated by non-epidemiological sources. Moreover, while some EBM hierarchies explicitly include non-epidemiological evidence, others (like the GRADE hierarchy in Figure 5.1B) do not (though non-epidemiological evidence is occasionally discussed in the GRADE literature); critical appraisal of evidence in EBM is oriented toward the critical appraisal of clinical epidemiology evidence.

Elsewhere, I discuss two competing frameworks for medical evidence, a "measurement framework" versus a "logic framework" or "argument framework" (Fuller et al. 2024). They are frameworks for identifying, evaluating, and synthesizing medical evidence. The measurement framework is that of EBM and includes its theory of what makes something evidence (the measurement theory), its methods for evaluating evidence (e.g., the hierarchies of evidence), and its methods for synthesizing evidence (mainly, meta-analysis). The argument framework is a promising alternative in need of further development (I will not carry out this development here); it includes its own theory of what makes something evidence (the argument theory), its own methods for evaluating evidence (which include those found in

disciplines studying reasoning, such as logic, statistics, and epistemology), and its own method for synthesizing evidence, which we can call "logical synthesis" or "argumentation synthesis." Argumentation synthesis integrates diverse lines of evidence supporting different premises in an argument through the force of inference. If the central icons representing the measurement framework are evidence hierarchies and forest plots (graphical representations of meta-analyses), then the main images of the argument framework are argument diagrams (Copi et al. 2014). Figure 5.2 displays a nodal diagram representing the Risk GP Model for treatment predictions, using the example of statin therapy.

We can illustrate the argument theory (and the argument framework) using the Risk GP argument scheme. Even though epidemiological measurements are the starting point in the Risk GP Model, the argument theory better serves the Risk GP argument scheme than the measurement theory of evidence because it shows how epidemiological measurements are relevant to prognostic and therapeutic prediction claims and creates room for other evidence that can confirm/disconfirm the scheme's empirical assumptions. Consider the hypothetical prediction claim "A statin will reduce the probability that Priya will have a cardiovascular event by 30% (in relative terms)," lowering the probability from, say, 10% to 7%. Let's say that Priya is a sixty-year-old patient in a physician's general practice with type 2 diabetes who has no cardiovascular disease, and the target population consists of all patients like Priya in the physician's practice. The first empirical premise in the Risk GP argument scheme claims that the effect size due to a statin in the pooled study population, measured using the relative risk reduction, is 30%.[31] Evidence for this claim ($e_1$ in Figure 5.2) is provided by the results of the USPSTF meta-analysis of RCTs (Chou et al. 2022), which measured a relative risk reduction of 30%, and thus (assuming the measurement is unbiased) well establishes the claim.

The next empirical premise is the sufficient similarity assumption. As I mentioned, USPSTF is relatively cautious about matching the target population for statins to the study population (Mangione et al. 2022), including in the target population only people who share risk factors with the trial participants, and excluding patients over age seventy-five from the

[31] Though we will see in Chapter 6 that the ARR is more informative for Priya and her physician, the USPSTF and ACC/AHA recommend estimating the ARR from the RRR measured in the trials along with an estimate of the baseline risk made using a risk calculator (Grundy et al. 2019; Mangione et al. 2022). In statin trials, the ARR varies according to baseline risk as well as demographic and clinical characteristics, while the RRR does not (Chou et al. 2022).

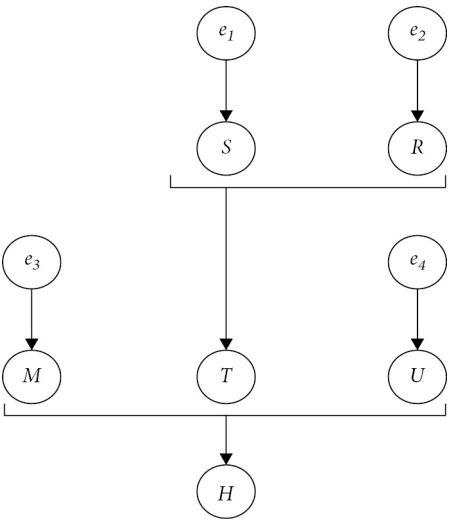



**Figure 5.2** Nodal diagram representing the Risk GP Model for a statin effectiveness prediction. Nodes represent propositions, arrows represent inferential support, brackets represent joint support. H: a statin will reduce the probability that the patient will have a cardiovascular event by 30% (in relative terms). T: the effect size in the target population will be 30% (RRR = 30%). S: the effect size in the study population is 30% (RRR = 30%). R: representativeness/sufficient similarity assumption. M: population membership assumption. U: unbiased population assumption. $e_i$: evidence.

target population because older patients were underrepresented in the trial populations (Chou et al. 2022). The more similar the target is to the study, the more likely it is to be similar in relevant ways. While I have not yet articulated the sufficient similarity assumption, in section 5.5 I provided examples of the kinds of causal-mechanistic details that are relevant.

An important intermediate stage in a statin's mechanism of action is a lowering of blood LDL-cholesterol (Borén et al. 2020), and statins have indeed been shown to reduce LDL-cholesterol in trials (Baigent et al. 2005). Thus the details included in the sufficient similarity assumption could include a reduced LDL-cholesterol, or the presence of other causes that interact with a statin prescription to lower LDL-cholesterol such as patient adherence. To extrapolate from the study population, the target population consisting of patients like Priya should resemble the study patients in these respects of similarity. To provide some evidence for the sufficient similarity

assumption ($e_2$ in Figure 5.2), Priya's doctor could check her LDL-cholesterol a few weeks after starting the statin to see if it has responded; if it's not practical to do so, they could, as a substitute, check with Priya to see if she is taking the statin as prescribed. If the results suggest reduced cholesterol, they support the sufficient similarity assumption and provide confirming evidence; if they do not, they undercut the assumption and provide disconfirming evidence. (If no information is available, the physician must make an assumption about whether the statin is having its intended effect on cholesterol.)

Consistent with my suggestion, ACC/AHA (Grundy et al. 2019) suggests checking LDL-cholesterol after starting statin treatment to assess adherence and response (and titrating statin dose to achieve at least a 30% to 50% reduction in LDL-cholesterol). USPSTF (Mangione et al. 2022) makes no such recommendation. USPSTF notes that most trials do not treat to a target LDL-cholesterol level and there are no trials that compare this strategy head to head with USPSTF's recommended fixed-dose strategy (while comparing the strategies across trials shows no differences in outcomes; Chou et al. 2022). However, though a dose titration strategy may not be shown to be superior, that does not mean that checking LDL-cholesterol response in practice is irrelevant because, as I mentioned, LDL-cholesterol is lowered by statins in RCTs and plausibly must be lowered to achieve any benefit (according to a statin's known mechanism of action), even if evidence to support a particular treatment target is lacking. In fact, in its statin guideline, USPSTF mentions that in one of the only trials assessing statins for older people, statins did not achieve a reduction in mortality. They note that among the problems with the trial was that there was little difference in LDL-cholesterol lowering between treatment and control groups, presumably because USPSTF thinks LDL-cholesterol lowering is relevant in assessing the effectiveness of statins.

If the sufficient similarity assumption has some support, then together with the first premise it provides some warrant for the intermediate conclusion that the relative risk reduction in the target population will be (approximately) 30%. While GRADE considers mechanisms in evaluating the indirectness or external validity of epidemiological evidence (Guyatt, Oxman, Kunz, et al. 2011), it uses information about mechanisms to evaluate epidemiological evidence rather than considering that information to be evidence on par with epidemiological study results. Moreover, the guidance that GRADE provides on evaluating evidence of mechanisms is sparse (Perillat and Mercuri 2022).

Next, evidence for the population membership assumption ($e_3$ in Figure 5.2), the assumption that Priya is a member of the target population, might consist partly in the results of a fasting blood sugar test supporting the diagnosis of type 2 diabetes and partly in the results of clinical epidemiology evidence demonstrating good specificity (positive accuracy) of positive test results for the diagnosis of diabetes. Evidence for the unbiased population assumption ($e_4$ in Figure 5.2), the assumption that assessment for statin therapy is not correlated with statin responsiveness, might consider indicators for statin adherence or nonadherence, which influences long-term treatment responsiveness. For instance, if Priya was assessed for statin treatment at a periodic diabetes health exam, is there any evidence to suggest that patients that regularly attend these exams are more likely to regularly take their medications (perhaps because they are more concerned about their health or have good health insurance that will cover these visits as well as statin treatment)? Together, all these empirical premises provide some warrant for the hypothesis that a moderate-intensity statin will lower the probability that Priya will have a cardiovascular event by 30% in relative terms. When combined with a prediction about Priya's baseline probability of a cardiovascular event within a wider argument, this hypothesis can help support a further conclusion about the extent to which a statin will lower the probability of a cardiovascular event in absolute terms.[32]

The more inclusive approach to evidence offered by the argument theory aligns better with the public health epidemiology tradition rather than the clinical epidemiology tradition (Fuller 2020c).[33] The classic public health epidemiology approach was on display during one of its formative moments, the discovery in the 1950s and 1960s that smoking cigarettes causes lung cancer, discussed in Chapter 3. The 1964 Surgeon General's Report for the U.S. Public Health Service considered a diversity of evidence in assessing the hypothesis that smoking cigarettes causes death from lung cancer, including non-experimental epidemiologic studies, clinical studies, and animal

[32] As this example demonstrates, the evidence needed for therapeutic predictions goes beyond evidence from "difference-making" clinical studies (such as clinical trials) and evidence of mechanisms, which are the two kinds of evidence integrated in the EBM+ approach (Clarke et al. 2014; Parkkinen et al. 2018).

[33] American epidemiologist Wade Hampton Frost (1941) suggested that the epidemiologist's job was not to focus narrowly on a single experiment but to make coherent a diverse body of evidence, writing, "Epidemiology is something more than the total of its established facts. It includes their orderly arrangement into chains of inference which extends more or less beyond the bonds of direct observation" (quoted in Parascandola 2011b, 246–7).

experiments. Abraham Lilienfeld, one of the public health epidemiologists contributing to the debate, said that "[t]he plausibility of the causal hypothesis is assessed, not in terms of the results of one particular study or a 'crucial experiment,' but in terms of the totality of available biologic evidence" (1959, 46). Another contributor, statistician Jerome Cornfield, suggested that there are no "hierarchies of evidence" to guide this assessment (Cornfield et al. 1959). Meanwhile, statistician R. A. Fisher, who continually questioned the conclusion that cigarettes cause lung cancer, argued that only a randomized trial provided "entirely unchallengeable conclusions" (1958b, 153). Fisher had developed the statistical theory of randomized experiments in the 1930s, and his approach to scientific inference generally evaluated a hypothesis using evidence from a single isolated statistical study (Parascandola 2005).[34]

I will not discuss all the virtues and vices of the measurement and argument theories. My suggestion here is that the argument theory can overcome some of the problems faced by the measurement theory. The main advantage of the argument theory of medical evidence over the theory that medical evidence consists in clinical measurements that I have been advertising so far is that it is better able to accommodate evidence for assumptions in our predictive inferences, evidence that is not easily understood as measuring something. The argument theory also offers a fruitful way of thinking about the results of clinical trials. Cartwright and Hardie (2012) construe positive RCT results, those showing a correlation between the intervention and the trial outcome, as establishing a premise in an "RCT Argument." When an RCT is unbiased, another premise in the RCT Argument, claiming that confounding causal factors are not associated with the intervention, is also satisfied. Together these premises help to warrant a causal conclusion about the study population.

An argument theory and argument framework for medical evidence are not necessarily opposed to the use of measurement tools in evaluating clinical epidemiology studies. The argument theory merely disagrees with the idea that what makes the results of clinical epidemiology studies medical evidence is that they measure clinical outcomes. We can still think about these results as resulting from measurements, and in so doing draw

[34] On the other hand, in the smoking controversy, Fisher conceded that an RCT could not be run to settle the question. So he sought evidence from twin studies to test his alternative hypothesis that a gene was a common cause of the association between smoking and lung cancer (Fisher 1958a). He thus reluctantly conceded the relevance of alternative sources of evidence in ruling out a rival hypothesis.

on measurement theory, measurement constructs, and measurement tools (such as risk-of-bias instruments) from clinical epidemiology to understand when and how these measurements are unbiased as well as to quantify statistical error. Measurement is one way of producing the data that serve as evidence for a medical hypothesis. Rather than hierarchies of evidence, the EBM hierarchies are better seen as hierarchies of measurement accuracy. Reconstruing them as hierarchies of measurement accuracy and combining them with the argument theory for evidence will not assuage the multiplicity of concerns that have been raised about the hierarchies. However, it will help to alleviate the two problems with evidence hierarchies introduced earlier by placing epidemiological measurement in its proper place.

The first problem is an exaggerated reliability of evidence near the top of the hierarchies, which is only as reliable for prognostic and therapeutic prediction as the predictive inferences that it supports. When critical appraisal of evidence focuses too narrowly on the unbiasedness of measurement, it neglects other determinants of evidential warrant that are also relevant to assessing the strength of evidence according to the argument theory, such as relevance (Cartwright 2009; Cartwright and Hardie 2012). Consequently, the strength of evidence provided by studies such as RCTs that are often less biased may nonetheless be overstated. On the argument theory, evidential relevance and strength are relative to an argument, and warrant for a hypothesis must be assessed by considering the overall argument and the empirical support for its various premises. The quality of the evidence for each premise must also be assessed. While clinical epidemiology offers tools and norms for evaluating the quality of clinical epidemiology evidence, evidence from other sciences might need their own principles. We need, for instance, principles for assessing the quality of mechanistic reasoning or evidence of mechanisms (Howick 2011; Parkkinen et al. 2018).

The second problem with the evidence hierarchies is that they provide no means of synthesizing diverse evidence. The argument theory does not nearly provide a comprehensive solution to the problems of integrating dissimilar kinds of medical evidence.[35] Nor need it reject the use of methods such as meta-analysis, the results of which can be used to establish a premise in a predictive inference, as we saw with the statin example. However, meta-analysis is narrowly applicable to the results of similar studies, such as a

[35] See Fletcher et al. (2019) for an overview of problems. Landes et al. (2018) propose an evidence synthesis framework for pharmacology that uses Bayesian nets (see also Stegenga and Menon 2017). Bayesian nets can be considered arguments that come with particular confirmation rules.

collection of RCTs. The argument theory provides a way of combining diverse evidence that cannot be pooled in a meta-analysis: diverse evidence is combined through argumentation synthesis when it supports different premises in the same argument. Diverse kinds of evidence are often seen as complementary, and rightly so on the argument theory when they support different premises, especially when those premises jointly support the argument's conclusion. The argument theory thus offers a defense of the value of diverse medical evidence that differs from a common robustness defense, which instead argues that evidence generated using more diverse methods is less sensitive to the assumptions or weaknesses of any one of these methods (Cartwright 2021). The defense also differs from the thesis that more diverse evidence is better able to eliminate competing hypotheses (Landes 2020).

## 5.8. Conclusion

An argument theory of medical evidence can help overcome several problems faced by EBM's measurement theory, including problems with its approach to evidence evaluation and medical prediction. The EBM hierarchies assume a measurement theory of medical evidence, in which evidence consists in measurements of clinical outcomes, especially risks. These are the kinds of outcomes that are readily measured in clinical epidemiology, the science of EBM. EBM and clinical epidemiology descend from an empiric tradition of medical prediction, which explains their measurement orientation. In contrast, I defended an argument theory of medical evidence, in which evidence is more broadly facts or data that establish premises in an argument for a medical hypothesis. The advantages of the argument theory over the measurement theory come into focus through an examination of the Risk Generalization-Particularization Model, EBM's standard model of prediction. In the Risk GP Model, epidemiological measurements are extrapolated to a target population (Generalization) and subsequently transformed into patient probabilities (Particularization). Unlike the measurement theory, the argument theory makes room for diverse evidence that can support all the model's assumptions. The argument theory also helps overcome two commonly discussed problems with evidence hierarchies that the measurement theory, by comparison, aggravates.

Over the next three chapters, I continue my analysis of evidence and inference in the new modern medicine by interpreting Particularization's several concepts of risk (in Chapter 6), by providing solutions to the problems of extrapolation or Generalization of effect sizes (in Chapter 7), and by defending a meta-evidentialist approach to the problem of biased medical evidence (in Chapter 8).

# 6
# Medical Interventions and Risk


## Summary

In grappling with the meaning of risk in medicine, two challenges confront us. First, how should we understand the effect of a preventive or curative medical intervention, which is often quantified using population-level effect measures such as the absolute risk reduction? Second, in what sense (if any) does epidemiological evidence reveal the individual "patient's risk" or their "reduction in risk" due to the intervention? I argue that preventive and curative medical interventions reduce the risk of medical conditions in a population compared to their omission, commonly by disrupting a mechanism for the condition, but that their effects on individuals vary according to the presence of mechanisms and causes that interact with the intervention to disrupt them. Thus, rather than inferring some fact about the patient from epidemiological effect measures—the patient's risk or their risk reduction—we should interpret effect measures as informing our uncertainty about the patient's outcome and their response to medical intervention.

In the "risky" new modern medicine, epidemiological evidence is seen as providing information about population-level as well as individual-level "risk" and "risk reduction." Preventive and curative medical interventions can be defined through their ability to reduce risk in a population compared to that same population if untreated, and epidemiological effect measures as quantifying this effect. Medical interventions commonly achieve their aggregate effect by disrupting disease mechanisms in individuals. However, treatment effect heterogeneity is blurred in the aggregate. That makes it tenuous to infer the patient's risk or their degree of risk reduction from the epidemiological evidence and effect measures. The patient's risk is a patient-level probability and is most naturally interpreted as a propensity, a physiological tendency toward the outcome. However, this propensity interpretation faces epistemological and practical problems. Thus I propose to reinterpret patient-level probabilities inferred from epidemiological evidence as degrees of belief, or credences. These credences capture epidemiological uncertainty,



*The New Modern Medicine*. Jonathan Fuller, Oxford University Press. © Oxford University Press 2025.
DOI: 10.1093/9780190066178.003.0007

or uncertainty about the distribution of individual outcomes in a population and where in the distribution the particular patient is located.

## 6.1. The Meaning of Medical Risk

Patients are not statistics. This statement is trite but often underappreciated in an era in which medical care is driven by evidence from epidemiological studies. The starting point in contemporary epidemiological medicine is a population. Using the Risk Generalization-Particularization Model from Chapter 5, we predict what will happen to a patient from the results of an epidemiological study. In Chapter 7, I dissect the first stage of this inference, Generalization of a treatment effect to a target population that includes the patient. I discuss the problems of extrapolation and the external validity of clinical trials. In this chapter, I probe the meaning of the information being extrapolated, which typically includes population and patient risks.

In epidemiology, risk is a technical term often referring to the proportion of a population with a health outcome. Risk is also a critical beam for contemporary scientific medicine. Many causes in medicine are risk factors for noncommunicable conditions like cancers or chronic diseases and are predictive of an increased risk of those conditions. Some chronic diseases, such as hypertension, type 2 diabetes, and obesity, are even themselves risk factors for other diseases. Risk is also an important concept for evidence-based medicine (EBM), given its origins in clinical epidemiology, a discipline consumed with quantifying risk. In fact, I argued in Chapter 5 that EBM views medical evidence as measurements of clinical outcomes, especially risks.

One of the dramatic shifts in the new modern medicine of today is a broadening of the central goals of medical intervention to include the prevention of chronic and noncommunicable diseases as well as their "management" or "control." This pivot in medical priorities was necessary because most chronic and many noncommunicable diseases are (presently) incurable, and thus the traditional curative goal of medicine (Broadbent 2019) is often elusive. That is not to say that cure isn't presently an important clinical aim, nor that prevention wasn't previously one. Prevention of disease through public health and hygiene was considered one of the essential duties of medicine in early twentieth-century medicine (Osler 1907) and was responsible for some of medicine's earliest and most significant victories against disease. However,

"preventive medicine" was mostly a term uttered in connection with efforts to stop the transmission of infectious disease, especially through population-level efforts such as mass vaccination, public water sanitation, and public education about germs.

Robert Aronowitz (2015) and Jeremy Greene (2007) tell of a new preventive medicine born in the second half of the century with the purpose of preventing noncommunicable diseases. Rather than populations, prevention was to be directed toward individuals at high risk of disease. The results of the Framingham Heart Study provided the initial impetus in the 1950s and 1960s by identifying individual-level risk factors for heart disease—high cholesterol, high blood pressure, high blood sugar—that could serve as targets for preventive medical intervention. In 1964, several Framingham investigators described this "[p]reventive medicine" as "largely the prevention of disease or disability through individual effort, both by the physician and the patient" (R. Dawber et al. 1964).

The peculiar new dialect of this preventive medicine is the language of risk, which leads Aronowitz (2015) to describe our current medical system as "risky medicine." Not only do the new preventive interventions target risk factors or "risk-based diseases" (Schwartz 2008) like hypertension, but the effectiveness of these interventions is couched in terms of their ability to "reduce risk" and is measured using measures such as the absolute risk reduction (ARR). It is not only primary prevention of disease (e.g., screening and treatment to prevent osteoporotic bone fractures) that is now "risky" (Viswanathan et al. 2018). "Management" of disease, the prevention of complications or progression of disease (e.g., using antidiabetic agents to prevent diabetic complications), is also framed in the language of risk reduction (American Diabetes Association 2021a). Disease surveillance such as routine monitoring in cancer survivors is guided by the risks of cancer recurrence and the risks of late effects of cancer treatment (Mulder et al. 2020). Even the goal of non-metastatic cancer treatment in current medical practice is less often understood as the "cure" of cancer but rather as the "prevention of cancer recurrence" (Waks and Winer 2019), which is similarly measured through a reduction in risk.

Risk-reducing interventions include lifestyle modifications: healthy behaviors like eating better and exercising more. However, risk-reducing medical interventions have also become among the most widely prescribed and used medical resources. These are most commonly medications but can include surgeries (such as breast- or colon-excising surgeries to prevent

familial cancers) and screening (such as breast and colon cancer screening to detect and treat cancer early). In the second half of the 1900s, the first classes of risk-reducing medications arrived, initiating a "therapeutics of risk reduction" (Greene 2007, 4). In the 1950s, Diuril (an antihypertensive) for hypertension and Orinase (an oral antidiabetic agent) for type 2 diabetes provided the prototype, followed later by Mevacor (a statin) for cholesterol in the 1980s. Each disease was initially defined and treated according to the presence of symptoms or signs, but after the asymptomatic were found to harbor a high risk of heart attacks and strokes as well as other complications—and after antihypertensives, oral antidiabetic agents, and statins were found to reduce the risk of these misfortunes—each class of medication became used primarily for treating risk in the asymptomatic (Greene 2007). A popular new class of drugs, the glucagon-like peptide-1 receptor agonists, is used to manage body weight and promote weight loss in people with obesity or who are overweight as well as to control blood sugar in people with type 2 diabetes. Recently one of these drugs, semaglutide, was also approved for some patients to reduce the risk of cardiovascular events (Food and Drug Administration 2024).

The historic conceptual shift toward risk management is most evident in the case of cholesterol-lowering therapy. It is natural to think that cholesterol-lowering drugs treat high cholesterol, or the disease known as "hypercholesterolemia." However, most patients receiving a cholesterol-lowering statin do not have high cholesterol above the upper limit of normal. Instead, they have a high risk of cardiovascular disease events like heart attacks and strokes. Statins are not primarily used to treat hypercholesterolemia. They are primarily used to treat cardiovascular risk. We could also say that they are used to prevent heart attacks and strokes, which is a less mysterious idea. However, considering that the population-level ARR attributed to statins is far less than 50%, either statins prevent a cardiovascular event in a minority of people or, more mysteriously, they lower an individual's risk of heart attack and stroke. Risk management thoroughly challenges our ordinary concepts of medical intervention.[1]

[1] In 2017, the American Heart Association (AHA) guidelines for the management of high blood pressure began following the cardiovascular risk paradigm, recommending for the first time that physicians target antihypertensive drug therapy to individuals at high risk of cardiovascular events and not just to those with the disease hypertension (Whelton et al. 2018). In 2023, the AHA released a consensus definition for a new condition, cardiovascular-kidney-metabolic (CKM) syndrome, which results from the interaction of cardiovascular disease, kidney disease, and metabolic diseases. One of the rationales for treating CKM syndrome as a diagnostic entity was that CKM confers a high risk of medical harms, especially cardiovascular events, and CKM syndrome could harmonize

There are thus two related challenges to understanding medical interventions and risks in contemporary medicine. First, what makes for a preventive (or curative) medical intervention in contemporary risky medicine? If the answer is that it reduces the risk of a medical condition in a population, then what does it mean to reduce risk in a population, and how are population-level risk and risk reduction related to outcomes in individuals, especially considering that outcomes and treatment effects often vary among individuals? Second, in inferring a patient's probability of the condition or the difference in their probability due to the intervention from this population-level evidence, how should we interpret those patient probabilities? It is common in contemporary risky medicine to say that the patient is at risk for the condition and that the intervention can lower or reduce their risk (by an amount quantified by effect measures like the ARR). It is natural to interpret these probabilities as facts about the patient. But is this interpretation warranted? If not, how should we interpret patient-level probabilities inferred from epidemiological evidence?[2]

I argue that preventive and curative medical interventions reduce the risk of medical conditions in a population compared to their counterfactual omission, commonly by disrupting an etiological or pathological mechanism for the condition, but that their effects on individuals vary according to the presence of mechanisms and causes that interact with the intervention to disrupt them. Thus, rather than inferring some fact about the patient from epidemiological effect measures—the patient's risk or their risk reduction—we should interpret effect measures as informing our uncertainty about the patient's outcome and their response to medical intervention. In section 6.2, I ascend from individuals to epidemiologic populations and analyze how preventive and curative medical interventions work at both levels. Then in section 6.3, I descend from the population back down to the individual to interpret the patient probabilities inferred from epidemiologic evidence,

disease management and risk reduction efforts across diseases (Ndumele et al. 2023). These changes continue a recent medical shift from the conceptual targeting of risk factors to the conceptual targeting of risk itself (Greene 2008).

[2] The new preventive medicine/risky medicine raises important questions about the expanded scope of medicine; overdiagnosis, overtreatment, and medicalization; the value of screening; the effectiveness of contemporary medical interventions; and the insidious influence of private industry in modern healthcare. While I will not directly tackle all these problems here, the conceptual analysis and revisions to our concepts that I will suggest could prove useful in confronting them.

rejecting a propensity interpretation in favor of a credence interpretation grounded in epidemiological uncertainty.

## 6.2. From Individuals to Populations

In epidemiological research, the effects of medical interventions are observed in populations. Preventive and curative medical interventions reduce the risk of a medical condition in a population, and effect measures quantify the amount of risk reduction, the "effect size." At the same time, the effect of a preventive or curative medical intervention in a population is the aggregate of its effect on individuals in that population and is determined by the distribution of disease mechanisms as well as causes that interact with the medical intervention to disrupt them. In this section, I define a preventive or curative medical intervention and discuss how its population-level effect is an aggregate of heterogeneous individual treatment effects.

The concept of a preventive or curative intervention I propose starts from how such interventions are studied in contemporary medical research. Preventive and curative medical interventions are identified in comparative epidemiological studies, studies such as randomized controlled trials (RCTs). Whether an intervention has preventive or curative effectiveness can be established by comparing a study group exposed to the intervention with an unexposed study group, one that receives no intervention (or a placebo/sham intervention meant to mimic the intervention).

In Chapter 5, our paradigmatic examples of risk-reducing medical interventions were statins. Statins are cholesterol-lowering drugs that are used to prevent cardiovascular events like heart attacks and strokes that result from atherosclerosis, or cholesterol buildup in the walls of arteries. Statins are recognized to be preventive interventions because they are associated with a lowered risk of cardiovascular events in meta-analyses of RCTs (Chou et al. 2022). We can represent this association using a graphical decision aid like the one from the Mayo Clinic in Figure 6.1, showing the relationship between statin use and ten-year heart attack risk.[3]

[3] Figures 6.1–6.3 show the effects of interventions in specific populations defined by certain risk factors. The effects are predicted to vary in other populations. In the absence of other conditions, current guidelines recommend considering a statin for populations with a higher untreated absolute risk of heart attack than displayed in Figure 6.1 because the number of heart attacks averted is predicted to be greater in higher-risk populations (Grundy et al. 2019; Mangione et al. 2022).

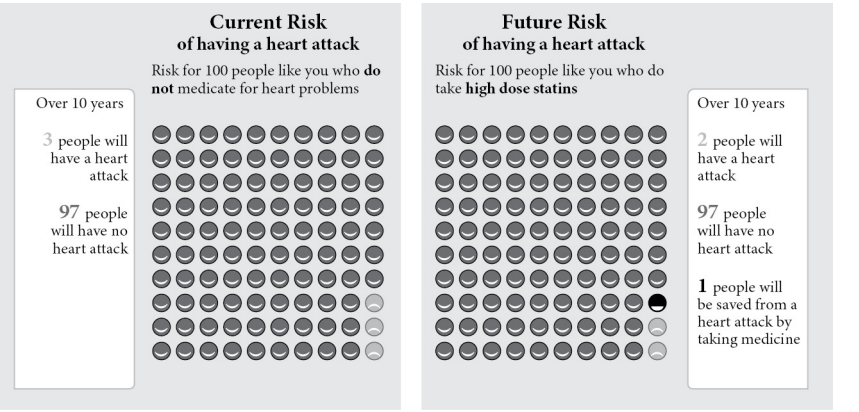



**Figure 6.1** Decision aid representing the effect of a statin on ten-year heart attack risk in a certain low-risk population. The two displays compare the population treated with a high-dose statin (right) to the population not treated with a statin (left).
Mayo Clinic (2024).

Decision aids are meant to help individual patients decide among interventions (or no intervention), yet they display the influence of interventions at the population level. This statin decision aid represents three population-level facts. On the left side, it represents the ten-year absolute risk of heart attack in an untreated low-risk population defined by certain risk factors. The absolute risk is the proportion of individuals in the population who will develop the event (in this case, a heart attack over ten years). The untreated absolute risk is 3/100, represented by three light gray people with heart attack out of 100 people. On the right side, the decision aid shows the absolute risk of heart attack in the population treated with a (high-dose) statin for ten years. The statin-treated absolute risk is 2/100, represented by two light gray people with heart attack out of 100 people. Finally, the graphic shows the effectiveness of a statin by comparing the scenarios on the right and left and showing that the absolute risk of heart attack is lower on the right by including a jet black person in the statin-treated population "saved from a heart attack by taking medicine."

Importantly, the population comparison in the decision aid is (implicitly) a counterfactual comparison, one in which two scenarios are being compared: population treatment versus population non-treatment over the same ten-year period. This is a "counterfactual comparison" in that at least

one of the two scenarios must not actually occur; it must be counter-to-the-fact. At least one scenario must not occur because it is impossible for the same population to be both treated and not treated with the same intervention at once. In the counterfactual approach to causal inference in contemporary epidemiology, causal effects are defined through a counterfactual comparison of this kind (Greenland and Robbins 1986). Similarly, we can define a preventive or curative intervention as an intervention that lowers the risk of an outcome in some actual treated population, where the risk is specifically lower compared to a counterfactual scenario in which that population is instead untreated.[4] According to the counterfactual approach to epidemiological causal inference, an unbiased treated-untreated comparison in a study such as an RCT approximates this counterfactual comparison.[5] This counterfactual account of a preventive or curative intervention is consistent with a natural counterfactual understanding of individual-level prevention and cure (Fuller 2022b).

We might, then, instead define a preventive intervention as one that prevents outcomes in some individuals, and a curative intervention as one that cures outcomes in some individuals. However, on a natural understanding of prevention and cure, prevention or cure only occurs when the outcome would have occurred otherwise. If this counterfactual "would have" is deterministic, where "would have" means that the probability of the outcome's occurrence in the counterfactual is 100%, then any medical outcomes that are chancy, or subject to chance, cannot be prevented or cured. Since some medical outcomes, such as cancer, are plausibly chancy because they arise partly through events like mutations that can be modeled stochastically (Armitage and Doll 1961), and because it is unclear whether medical outcomes are generally deterministic, it is plausible that were we

[4] More precisely, an intervention is a preventive or curative intervention if and only if for some actual population the intervention occurred and the probability of the outcome in the population was p, and there occurred an X such that (1) there is a causal interaction between the intervention and either X or the process due to X, and (2) if the intervention had not occurred, then X or the process due to X would have fixed a probability of the outcome greater than p (Fuller 2022b). For a medical intervention, the outcome is a health outcome, X is some cause, and the process is an etiological or constitutive disease mechanism with which the medical intervention interacts. This characterization of a *medical* intervention is arguably overinclusive, including, for instance, economic interventions that reduce the risk of poverty-related diseases by interacting with poverty-related social mechanisms.

[5] The counterfactual comparison must be a treated-untreated comparison rather than a comparison among different interventions because in a counterfactual comparison of different but equally effective interventions, there may be no net risk lowering. A placebo- or sham-controlled trial should be considered an approximation to a treated-untreated counterfactual comparison only when the placebo/sham is not itself a preventive or curative medical intervention with respect to the outcome.

to define a preventive or curative intervention as one that prevents or cures outcomes, it would have the undesirable implication that some—possibly many—of our stock examples of medical interventions may not count as preventive or curative interventions. My population-level concept of a preventive or curative intervention is agnostic as to how preventive and curative interventions bring about population-level risk-lowering, whether by preventing/curing disease deterministically in individuals comprising the population or by lowering risk in these individuals (or by doing both). It accords with epidemiological and medical practice by identifying preventive and curative interventions as ones that show population-level risk-lowering in an epidemiological comparison.

Although it is more common to describe preventive interventions rather than curative interventions as "risk-lowering," that may reflect the increased prominence of medical prevention in contemporary epidemiological medicine. Even interventions like cancer treatments that may well eliminate cancer in some people are often described as preventing cancer recurrence (Waks and Winer 2019), cancer progression, or death from cancer because their effectiveness is established through comparisons in which the risk of future outcomes is reduced. However, there are other contemporary examples of interventions, such as newer antimicrobial regimens, that *are* recognized by medicine as curing or "eradicating" disease, such as medication regimens for treating *H. pylori* infection (Zamani et al. 2022). While preventive interventions reduce the risk of the occurrence of disease, curative interventions reduce the risk of the persistence of disease.

Preventive and curative *medical* interventions often work by disrupting a mechanism[6] or eliminating some pathology. Preventive medical interventions often disrupt an etiological mechanism that generates a disease or medical condition, while curative medical interventions often disrupt a constitutive or pathological mechanism that sustains a disease or medical condition. Examples of pathological mechanisms that sustain a disease are mechanisms of pathogen replication and survival, as well as mechanisms of cancer cell replication and survival; examples of curative medical interventions that disrupt them are antibiotics/antivirals and cancer chemotherapies, respectively. Other times, a disease is "sustained" by some pathology consisting of structural lesions in the body, such as osteoarthritis

[6] The basic idea behind mechanism disruption is that the intervention interrupts the causal chain from upstream factors to downstream factors (Fuller 2022b). See also Steel (2008).

or gross tumors. And curative medical interventions sometimes work by eliminating or removing this pathology from the body, such as through joint replacement or through the surgical removal of tumors.

Examples of etiological mechanisms that generate a disease or medical condition include mechanisms of pathogen transmission and infection, atherosclerosis, mechanisms of atherosclerotic plaque rupture and thrombosis (causing heart attack or stroke), and various mechanisms by which cancer causes death; examples of preventive medical interventions that disrupt them are vaccines, statins, carotid endarterectomy, and adjuvant cancer treatments, respectively. Vaccines disrupt mechanisms of pathogen infection by providing immune memory that is triggered by the pathogen and kills it. Statins lower circulating LDL-cholesterol, disrupting atherosclerotic plaque buildup in arteries. Carotid endarterectomy disrupts mechanisms of plaque rupture in carotid arteries that supply the brain with blood by surgically removing the plaque. And the adjuvant cancer treatment trastuzumab, a monoclonal antibody, binds to breast cancer cells and kills them before they have a chance to kill the patient.

In summary, preventive and curative medical interventions reduce the risk of a medical condition in a population compared to their counterfactual omission, commonly by disrupting an etiological or pathological mechanism for the condition. However, they rarely do so by disrupting a disease mechanism in everyone. Not everyone has a disease mechanism to begin with. And a medical intervention will generally not disrupt it in everyone. (Medical interventions are rarely 100% effective.) This heterogeneity in outcomes among individuals complicates interpreting risks and risk reductions for individuals, as we will see in section 6.3. (It also complicates extrapolating from study populations to target populations, as we will see in Chapter 7.)

In Chapter 3, we encountered epidemiology's sufficient-component cause model (Rothman 1976), a multifactorial theory of the etiology of medical conditions. On this theory, causes, including medical interventions, are not sufficient on their own to effect medical outcomes; they are merely component causes within sufficient causes of those medical outcomes. The other components with which a medical intervention causally interacts within its sufficient causes are its "cofactors" (Greenland and Brumback 2002). These will include "disruption cofactors" that interact with the intervention to help it disrupt its targeted mechanism. For all interventions, disruption cofactors include patient adherence to the intervention. For drugs, they also include

factors involved in mechanisms of absorption, distribution, metabolism, and elimination of the drug, such as metabolic enzymes that maintain adequate concentrations of the active form of the drug in the body (Lai et al. 2022). For vaccines, disruption cofactors include immunologic factors involved in acquiring immune memory, which people with certain immune-compromising conditions may lack (Ljungman 2012). Finally, for targeted treatments, those that bind selectively to a biomarker, disruption cofactors include the presence of that biomarker. For instance, trastuzumab binds selectively to the HER2 cell surface receptor, which is overexpressed by cancer cells in some cases of breast cancer (Vu and Claret 2012).

As seen with heart attacks in Figure 6.1, there is generally heterogeneity in individual outcomes within an untreated population. Not everyone has a heart attack or dies from cancer over a ten-year period (thankfully). In general, there are two potential reasons for this heterogeneity in untreated outcome. It could be that not everyone has a mechanism for the medical condition. Atherosclerosis does not produce an unstable coronary artery plaque in everyone. It could also be that the outcome is governed by chance: in those with the mechanism, it produces the outcome with a probability of less than 100%. At one extreme, perhaps everyone in the untreated population in Figure 6.1 has a 3% individual probability of a heart attack. In section 6.3, I will describe this latter possibility by saying that people with the mechanism have an individual "propensity" toward the outcome. It could even be that different mechanism types or subtypes are associated with a propensity of a different strength, such that different mechanisms produce the outcome with a unique probability. The absolute risk in the untreated population, the frequency or probability of the outcome in the population, is determined by the distribution of mechanisms in the population—how many people have a complete undisrupted mechanism (and possibly also what type of mechanism they have). Differences in the absolute risk among untreated populations are common for outcomes like heart attack and cancer and are produced by differences in the population distributions of disease mechanisms.

There are also generally differences in individual treatment effects within a treated population (Kravitz et al. 2004; Angus and Chang 2021). A statin must have an effect in some individuals to produce a difference in absolute risk of heart attack, treated versus untreated, in Figure 6.1. But it must not successfully disrupt all heart attack mechanisms, or else there would be no heart attacks in the treated population. It could be that disruption cofactors

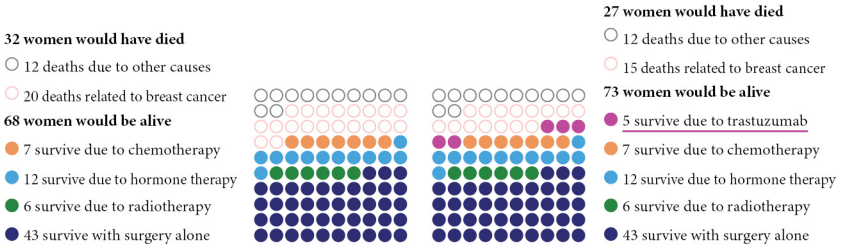



**Figure 6.2** Decision aid representing deaths in certain women with invasive breast cancer at fifteen years following breast cancer surgery. The two displays compare the population treated with trastuzumab and other standard interventions (right) to the population treated with other standard interventions alone (left).
National Health Service (2024).

are not distributed equally among those with clinically important atherosclerosis, so that atherosclerosis continues to build toward an unstable arterial plaque in some of them despite the statin. Or it could be that atherosclerotic plaque is already established in some people and that a statin does not disrupt all mechanisms by which established plaque results in heart attack (Borén et al. 2020).

If atherosclerotic plaque is established in everyone, then it's still possible that the treatment has an equal effect on everyone, lowering everyone's 3% individual probability of heart attack by 1% through disrupting further atherosclerosis, while leaving them each with a 2% probability of heart attack due to previously accumulated plaque. At the other extreme, it could be that the statin lowers the individual probability of heart attack from 100% to 0% for one person, while two individuals remain at 100% risk of heart attack due to established plaque.[7]

Figure 6.2 visualizes the contribution of mechanism heterogeneity to individual treatment effect heterogeneity using another example. It displays the risk of death in certain women treated after surgery for invasive breast cancer with trastuzumab and other standard interventions (right display) versus other standard interventions alone (left display). There are two kinds of mechanisms responsible for death in both scenarios: cancer mechanisms resulting in death from breast cancer, and non-cancer mechanisms

[7] As Stephen Senn (2004) notes, traditional clinical trials cannot distinguish between these two kinds of explanations for outcome variability.

resulting in death from other causes. Comparing the trastuzumab-treated and trastuzumab-untreated populations, trastuzumab predictably makes a difference only to the number of deaths due to cancer mechanisms because those are the mechanisms it targets and potentially disrupts. Whenever the outcome is one produced by diverse mechanisms, including mechanisms not targeted by the intervention (as with all-cause mortality), we should expect individual treatment effect heterogeneity. Another common example is the outcome "flu-like illness," which is produced by multiple different pathogen-specific etiological mechanisms. The influenza vaccine targets only a subset of them and thus will have potentially substantial individual treatment effect heterogeneity. In Chapter 4, I argued that many chronic diseases and mental disorders either are associated with known nonspecific pathology or may have unknown nonspecific pathology because they are classified according to nonspecific manifestations. Because some of this pathology may not be targeted by medical and psychiatric medications, those medications could have substantial individual treatment effect heterogeneity.

A final source of individual treatment effect heterogeneity is paradoxical harms, where the intervention helps cause the outcome in some individuals that it helps avert in others. Consider Figure 6.3, representing the effect of carotid endarterectomy on the frequency of stroke among those who'd experienced a previous stroke event. Unlike the previous two graphics, this one represents the effects of the intervention, and it separates benefit from harm. According to the authors, there are some individuals who "benefit" from

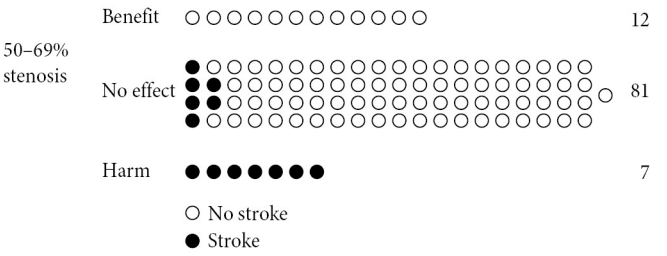



**Figure 6.3** The effect of carotid endarterectomy on five-year risk of stroke among those with 50–69% symptomatic carotid artery stenosis. Harms are operative strokes or operative death; all other outcomes refer to the occurrence or non-occurrence of ipsilateral stroke.
Rothwell et al. (2005, 257).

endarterectomy by having a stroke prevented. There are others who experience "no effect" and who won't have a stroke with or without surgery. There are still others under "No effect" who *will* have a stroke with or without surgery. And there is a final subpopulation of people for whom endarterectomy will cause an operative stroke (or, rarely, death), represented by people under "Harm." Endarterectomy displays all manner of individual outcome heterogeneity within a population. (In Chapter 7, we will see that it also displays aggregate outcome heterogeneity between populations.)

In the statin and trastuzumab graphics, individual benefits were not distinguished from individual harms; only the net number of outcomes averted was represented. It remains open from the data in Figures 6.1 and 6.2 whether those interventions are sometimes responsible for paradoxical harms, buried within the heart attacks or deaths in the treated group. Because operative strokes occur soon after and can be attributed to the surgery, paradoxical harms can often be identified for endarterectomy. Other examples of identifiable paradoxical harms occur when interventions that help avert deaths in some help produce deaths in others, as in other fatal surgical complications and fatal drug reactions in which there is a deadly mechanism attributable to the intervention. In another example, digoxin can both prevent hospitalization due to heart failure and cause hospitalization due to other conditions (Upshaw et al. 2018).

In summary, heterogeneity in individual outcomes within a population has several sources: differences in who has a disease mechanism targeted by the intervention, differences in who has a disease mechanism not targeted by the intervention, differences in who has disruption cofactors, and differences in who has harm cofactors. These differences aggregate to produce population risks and risk reductions. The untreated absolute risk is influenced by the frequency of disease mechanisms. Meanwhile, the ARR and relative risk reduction (RRR) can be influenced by the frequency of some or all of these causal variables, depending on the intervention, outcome, and population. (In the appendix, I model the influence of these variables on the absolute risks, ARR, and RRR under various assumptions.) The ARR is calculated by taking the difference in absolute risk between treated and untreated populations, while the RRR is calculated by dividing the ARR by the absolute risk in the untreated population. When the comparison between treatment and no treatment is a counterfactual comparison, then we can call the ARR and RRR the "causal ARR" and the "causal RRR," respectively, as measures of the effect of the intervention.

Alex Broadbent (2013) refers to the challenge of determining what causal information is conveyed by epidemiological measures like the causal ARR the "causal interpretation problem." His solution recognizes that effect measures aggregate singular causal facts because they are population measures and those populations are composed of individuals, each of whom is individually exposed or unexposed and has some individual response or nonresponse. Broadbent's solution relies on the apparatus of contrastive causal explanation: an association between two epidemiological groups is attributable to an exposure when that exposure explains the magnitude of association because it causes at least as many cases of the outcome in the exposed group as there are excess outcomes in that group. Then an effect measure can be causally interpreted as quantifying this singular causation. Broadbent says that the exposure must cause *at least* as many cases as there are excess outcomes in the exposed group to allow for the possibility that the exposure causes some cases that would have happened anyway, as when smoking causes lung cancer in those who are also exposed to a sufficient dose of a second carcinogen (resulting in causal preemption or causal overdetermination; Lewis 1973), and there are a similar number of nonsmokers who are exposed to a sufficient dose of the second carcinogen. These cases are not picked up by the effect measure, as a measure of the net effect calculated by comparing groups.[8]

Broadbent has in mind exposures that cause health outcomes rather than preventive or curative interventions, and thus his interpretation applies best to measures of positive effect, including measures of the "attributability" of excess outcomes to an exposure (Broadbent 2013). I am interested here in measures like the causal ARR and causal RRR that measure the effectiveness of preventive and curative interventions. Still, one could interpret the causal ARR and causal RRR along the lines of Broadbent's solution by saying that the intervention *prevents or cures* at least as many cases as there are fewer outcomes in the treated group. But to allow for chancy outcomes, I defined a preventive or curative intervention as one that reduces the risk in the treated population. This occurs when the intervention prevents, cures, *or reduces the risk of* the outcome for some members in the treated population, at least as many as needed to explain the amount of aggregate reduction in risk in

[8] Specifying that the exposure must cause *at least* as many cases as there are excess outcomes also allows for the possibility of paradoxical harms of the exposure that "cancel out" some of the cases caused by the exposure, as well as the possibility that the exposure exerts a causal influence on more people but the outcome is chancy and doesn't always occur under exposure.

the treated population. Because prevention and cure have a natural counterfactual connotation (Fuller 2022b), we can give an added counterfactual interpretation to measures of preventive or curative effect like the causal ARR, where the treated and untreated populations in which they are defined are the same population, treated versus untreated. Then measures of association like the ARR used in actual epidemiological studies provide an accurate estimate of the causal ARR when the study is unbiased.[9] This interpretation allows that effect measures are aggregating the effects of preventive and curative interventions on the individuals comprising the population.

Because epidemiological effect measures are aggregate measures, they do not on their own warrant causal conclusions about given individuals. From the causal ARR or causal RRR in the population, we cannot straightforwardly infer the effect on an individual, including the degree of "individual risk reduction." For instance, the absolute reduction in risk of stroke with carotid endarterectomy in Figure 6.3 is 5%, while the relative reduction in risk of stroke is 28%. But this clearly cannot mean that endarterectomy reduces the individual risk of stroke by 5% in absolute terms or by roughly a quarter in relative terms (from 18% to 13%) in everyone because endarterectomy causes operative stroke in some. In Figure 6.2, the absolute reduction in risk of death with trastuzumab is 5% and the relative reduction in risk of death is 16%. But trastuzumab must not achieve these amounts of risk reduction in everyone because it does not target mechanisms resulting in death from other causes. Finally, in Figure 6.1, the absolute reduction in risk of heart attack with a statin is 1%, while the relative reduction in risk of heart attack is 33%. But rather than lowering everyone's individual risk from 3% to 2%, a statin could, at the other extreme, prevent a heart attack in exactly one individual.

To recap, preventive and curative medical interventions reduce the risk of a medical condition in a population compared to their counterfactual omission, commonly by disrupting an etiological or pathological mechanism for the condition in individuals within the population. Their effects on individuals are often heterogeneous and are influenced by whether individuals have a mechanism for the condition as well as whether they have cofactors that interact with the intervention to disrupt a mechanism or

[9] Broadbent (2013) discusses several challenges with counterfactual interpretations of effect measures, at least some of which can be overcome by clearly specifying the counterfactual comparison and specifying it in such a way that an unbiased study comparison approximates it.

produce a paradoxical harm. Effect measures quantify the aggregate of these individual treatment effects, blending individual effects together to create a composite picture, the effect size. That makes the meaning and relevance of population risks and effect measures perplexing when we descend back down to the individual.

## 6.3. From Populations to Individuals

In the previous section, I interpreted population measures of an intervention's effect, especially the causal ARR and causal RRR, which quantify changes in the population's absolute risk of a medical outcome due to a preventive or curative medical intervention. Ultimately, these measures inform individual patient care, in which the population risks are transformed into "the patient's risk" and the risk reduction the patient can expect to receive from the intervention. Does the received interpretation of patient-level probabilities inferred from epidemiological evidence withstand scrutiny? If not, how should we reinterpret patient-level probabilities in epidemiological medicine? In section 6.3.1, I examine the concept of the patient's risk and reject interpreting patient probabilities inferred from epidemiological evidence as propensities. In section 6.3.2, I defend an alternative credence interpretation in which patient probabilities represent epidemiological uncertainty.

### 6.3.1 A Propensity Interpretation of the Patient's Risk

In providing evidence for medicine, epidemiological research measures risks. However, there are at least two concepts of risk operative in epidemiological medicine (Sniderman et al. 2015): absolute risks (and their effect size derivatives) and the patient's risk (and reductions in their risk). Epidemiological evidence directly measures absolute risks and absolute risk reductions. Because the patient's risk is inferred from the absolute risk and the change in the patient's risk is inferred from the absolute risk reduction, epidemiological evidence is also sometimes taken to measure the patient's risk or the change in the patient's risk. As we saw in Chapter 5, this leap requires inferences and assumptions. But what is this thing—the patient's risk—that medicine expends so much evidence and effort in trying to measure? And

should we interpret patient-level probabilities inferred from epidemiological evidence as the patient's risk and their reduction in risk?

The field of risk communication sprang up to study empirically how people (especially patients and physicians) understand risk and to enhance discussions through tools like graphical decision aids (Bonner et al. 2021)—the psychology of risk. Yet a philosophical interpretation of the patient's risk is lacking. In Chapter 5, I suggested that the patient's risk was an individual-level, single-case probability. But this answer only passes off the problem because it is not obvious how we should interpret single-case probabilities in medicine. Steven Goodman, a clinical epidemiologist, writes, "The concept of risk has undoubtedly become an enormously useful and essential part of medical science, but evidence-based medicine must grapple with the fact that in the individual patient, one of its pillars—the concept of probability—is ambiguous and elusive" (1999, 606).[10] Effective risk communication hinges on the problem of risk interpretation, with which philosophers can be of some help. Thankfully, philosophy of science has ready-to-order interpretations of probability for us to choose from.

In the original article in which the Evidence-Based Medicine Working Group announced its new "paradigm" to the medical community, they provided a vignette titled "The Way of the Future," in which a resident does a literature search to discover how likely it is that their patient will have a recurrent seizure. They find a study showing "the patient risk" of recurrence and communicate this information to the patient (EBM Working Group 1992). That future has arrived. Prognosis now involves individualized risk prediction. About her decision to have a double mastectomy, Angelina Jolie (2013) writes, "I carry a faulty gene, BRCA1, which sharply increases my risk of developing breast cancer and ovarian cancer. My doctors estimated that I had an 87 percent risk of breast cancer and a 50 percent risk of ovarian cancer, although the risk is different in the case of each woman."[11] Similarly, in evidence-based guidelines on the assessment of cardiovascular risk, the authors write, "It is important to note that risk estimation is based on group

[10] Similarly, Henrik Wulff, another clinical epidemiologist, suggests, "There is also the problem of clinical epidemiology being concerned with groups. I have yet to find a clinical epidemiologist who openly acknowledges the ambiguity of the probability concept. They see in a scientific study that 90 percent of patients were cured, and then they take it for granted that there is a 90 percent chance that Mr. Brown will be cured. But that is an extremely interesting problem. What is a probability of one person being cured?" (quoted in Daly 2005, 108).

[11] I borrow this example from Philip Dawid (2017), who surveys concepts of "individual risk" across diverse contexts. Here I am concerned only with the concept of individual risk operative in medicine.

averages, which are then applied to individual patients in practice" (Goff et al. 2014, S53).

There is good reason to suspect that, conceptually, the patient's risk is different from the absolute risk, the finite frequency of the outcome in the population. For starters, the patient's risk is about an individual, not a population. As we read in Jolie's quote, it can vary between two people. Moreover, we can alter the patient's risk by providing the patient (as opposed to the whole population) with a treatment. Jolie (2013) writes, "Once I knew that this was my reality, I decided to be proactive and to minimize the risk as much as I could. I made a decision to have a preventive double mastectomy." Similarly, according to evidence-based guidelines for cholesterol treatment, "statin therapy is recommended for individuals at increased ASCVD [atherosclerotic cardiovascular disease] risk who are most likely to experience a net benefit in terms of the potential for ASCVD risk reduction" (Stone et al. 2014, S7). The notes accompanying the statin decision aid in Figure 6.1 include the following suggested medical documentation: "We estimated the patient's 10-year risk of atherosclerotic events at 3% and discussed how this risk could be reduced with the use of statins to 2%." They also include the following note to the patient user about other benefits: "The use of statins reduces your stroke risk by about one fifth" (Mayo Clinic 2024).

The patient's risk is discussed as if it were a single case, individual patient probability. Probabilities are typically considered to be closely linked to finite frequencies like the absolute risk, yet distinct. The value of the patient's risk is taken to equal the absolute risk measured in a cohort study, and the magnitude of the change in the patient's risk due to treatment is taken to equal the ARR or RRR measured in an RCT or observational study. However, we saw in Chapter 5 that the patient's risk, as an individual probability, must be *inferred from* the absolute risk, and the change in the patient's risk due to treatment must be *inferred from* the ARR or RRR. How, then, does medicine interpret "the patient's risk," if not as the absolute risk or risk reduction in a population? For that matter, how can we interpret any probability? There are several ways of distinguishing the kinds of interpretations on offer, including subjective versus objective interpretations (Gillies 2000) and belief-type versus frequency-type interpretations (Hacking 2001). Both these dichotomies separate epistemic conceptions of probability (where probability refers to an epistemic notion like degrees of belief or the bearing of evidence on a hypothesis) from ontic conceptions of probability (where probability refers to something independent of our beliefs or evidence like frequencies or propensities).

I will first consider whether the patient's risk is likely an epistemic concept. On this view, the patient's risk might be a certain degree of belief the physician has (or should have) that the patient will develop the outcome. Or it might be the logical support that the hypothesis "The patient will develop the outcome" receives. However, a problem arises when we try to understand what physicians and patients mean when they say that a statin lowers a patient's risk of heart attack or that a double mastectomy lowers their risk of breast cancer. Seemingly, it means that the treatment causes these risks to be lower. On any realist theory of causation, causes exert a real effect in the world outside our minds. Sure, my belief that you will have a heart attack might be diminished if you take the medication, but the effect to which we are referring when we say that the intervention lowers your risk is not the effect the intervention has on my beliefs.

The patient's risk is more likely an ontic concept, locating probabilities in the world. One such interpretation to consider is frequentism, in which probabilities are either finite frequencies (e.g., 3% of people in the low-risk target population will have a heart attack) or "long-run" frequencies. Yet the finite frequency *just is* the absolute risk, a feature of the population rather than the individual. Meanwhile, the long-run or limiting frequency would be equivalent to the proportion of heart attacks if we drew a patient randomly from the target population and then replaced them, ad infinitum. That interpretation would only expand the population and our problem. Either variety of frequentism is prone to the notorious difficulty (if not outright futility) of making sense of a frequentist probability for a single case. Frequencies are categorically different from single cases; in fact, they are composed of a collection of single cases. Thus frequentism is not a good candidate for medicine's singularist concept of the patient's risk. As three epidemiologists put it, frequencies "cannot tell us anything about a particular case and therefore would leave a physician presented with a single patient in a quandary" (Djulbegovic et al. 2011, 310).

Among common ontic interpretations, the most plausible for the patient's risk is a propensity, a physical property of the single case that is disposed toward the outcome. Propensities are realizable properties, just like the dispositions discussed in Chapter 4. In that chapter, I analyzed chronic diseases and mental disorders as bodily dispositions toward clinical manifestations. Because I will ultimately endorse a different interpretation of patient probabilities provided by epidemiological evidence, I won't provide a detailed metaphysical analysis of patient propensities nor commit to any

particular philosophical account of propensity. (See Gillies 2000 and Suarez 2020 for historically oriented overviews of philosophical accounts.)

Like dispositions, propensities are characterized by their manifestations. Like bodily dispositions, bodily propensities may depend for their existence on a categorical base the way that a radioactive atom's propensity to decay depends on its categorical atomic structure. But unlike "sure-fire dispositions" that manifest their outcome reliably, propensities manifest their outcome in a chancy manner. If heart attacks are governed by propensities, then were the same pathological mechanism that generated a patient's heart attack to repeat, the second time a heart attack might not result. Consistent with the way that guidelines and patients speak about it, the patient's risk is most readily understood as their individual propensity toward some clinical outcome, and a lowering of the patient's risk due to an intervention is most readily understood as a change in that propensity. The value of the patient's risk would then represent the strength of that propensity or the frequency of the outcome that the propensity would manifest if its associated mechanism were to repeat indefinitely.

Many philosophers are quite comfortable ascribing propensities to decaying radioactive atoms or coin-tossings. Setting aside any doubts about the existence of propensities in these paradigm settings, are clinical outcomes chancy in the same way that these other outcomes are often thought to be? Some diseases depend on processes that are often thought to be chancy and are thus good candidates for medical outcomes that are governed by propensities. Recall that cancer results partly from mutations. Mutation is often modeled as a stochastic event (Armitage and Doll 1961). Carcinogens might raise the probability of a cancer mutation in an individual, but whether a cancer mutation results might depend partly on the play of chance, and thus the occurrence of cancer may be governed by propensities.

One problem with understanding patient risks as *intrinsic* bodily propensities is that the occurrence of a clinical outcome over a given period might depend partly on extrinsic causes during that period. The occurrence of cancers might depend on exposure to carcinogens, the occurrence of infection will depend on exposure to pathogens, and the occurrence of most diseases will depend on exposure to interventions. The probability that an individual will develop a clinical outcome is equal to the probability of the outcome given certain exposures as well as the probability of those exposures. Thus clinical propensities governing outcomes from cancer to infection would most likely be propensities of the world rather than propensities intrinsic to the patient.

For other medical outcomes, such as a bone fracture after a fall, it is less obvious that the mechanism leading to the fracture in a given case is chancy. (Perhaps the individual was shoved; perhaps a stochastic gust of wind knocked them off balance.) Clinical outcomes might only sometimes be governed by propensities.[12] Whether patient risks-as-propensities generally govern the diverse range of clinical outcomes encountered by doctors and their patients should be considered unsettled.

Thankfully, I do not need to settle the matter. My ultimate concern is how best to interpret the patient-level probabilities inferred from epidemiological research. Even if patient risks-as-propensities exist and generally govern clinical outcomes, we should not interpret the patient probabilities derived from epidemiological evidence as propensities. My arguments against interpreting them this way are mainly epistemic and pragmatic. I will remain agnostic as to whether patient propensities really do govern clinical outcomes. The problem is that the patient-level probabilities inferred from population-level absolute risks and risk reductions are not necessarily a good estimate for patient propensities. Moreover, misestimating the value of patient propensities could have real and serious consequences for patient care.

First, if we assume for the sake of argument that there are patient propensities, the aggregate propensity or the change in aggregate propensity in an epidemiological population may be an unreliable estimate of patient propensities or changes in propensity, respectively. As a measure of propensity, the absolute risk would first and foremost be a measure of average propensity in a population. The value of patient propensities could vary, or it could be similar among individuals. Recall from section 6.2 that medical conditions are produced by etiological or pathological mechanisms. The value of patient propensities could vary by the kind of mechanism present. (Atherosclerosis might on average be associated with a low propensity toward heart attack, but respiratory failure depriving the heart of oxygen might on average be associated with a high propensity toward heart attack.) The value of patient propensities could also vary among individuals with the

[12] DesAutels (2015) argues that many mechanisms in the life sciences are stochastic and that the probabilities attached to these stochastic mechanisms are best given a propensity interpretation. However, none of the mechanisms DesAutels discusses is a mechanism for clinical outcomes. The viability of a propensity interpretation of clinical outcomes need not hang on whether the generative mechanisms are all genuinely indeterministic. Probabilities associated with repeatable conditions such as a coin-tossing setup are sometimes thought to be *deterministic* propensities. Abrams (2017) notes that many mechanisms in the life sciences are recurrent, operating repeatedly. Thus they are good candidates for mechanisms associated with propensities, even if it turns out that they are not genuinely indeterministic.

same kind of mechanism. (Severe atherosclerosis might provide a greater propensity toward heart attack compared to minimal atherosclerosis.) And some individuals might not have a propensity at all if they lack a mechanism for the condition. By aggregating individual patient outcomes, the absolute risk in the untreated population washes out this variation.

Also recall from section 6.2 that preventive and curative medical interventions might have heterogeneous treatment effects for individuals in the population, and thus an uneven effect on patient propensities. Preventive and curative medical interventions work by disrupting the mechanisms described above. They won't affect individuals who lack a mechanism. Nor will they work on individuals who lack the disruption cofactors needed for the intervention to disrupt mechanisms. (Drugs will not work the same in patients who do not adhere to them or have drug-metabolizing enzyme variants.) In some individuals, they may even produce paradoxical harms (as with many adverse drug reactions). And when they do work, the amount of individual risk reduction may depend on the value of the initial propensity, as well as the presence of other mechanisms for the same outcome. (An individual might be saved from respiratory failure but still have atherosclerosis and remain at risk for heart attack.) By aggregating individual patient outcomes, the causal ARR and RRR wash out variation.

That said, the mere presence of variability in mechanisms or cofactors among individuals does not invalidate the inference from the aggregate to the individual, so long as we do not require that this inference is infallible and that the aggregate is a perfect estimate for all individuals. Ultimately, when inferring an individual propensity or change in propensity from the aggregate value, the question is whether the aggregate is a good estimate for many or most individuals. Whether or not an aggregate propensity (or change in propensity) is a good estimate for many individuals depends on the distribution of individual propensities (or changes in propensity) in the population. Observe Figure 6.4. If the population's distribution is narrow, normal, and unimodal (graph A), then the true value for most individuals will cluster around the mean, so that the aggregate will usually be a pretty good estimate for the individual. If the distribution is wide (graph B), skewed (graph C), or multimodal (graph D), then many individuals may cluster away from the mean, and the aggregate will often be a poor estimate for the individual. As an illustration, if the clinical outcome is *not* governed by chance, then individual probabilities and reductions in probability will have only two values, 0% or 100%, which means that the absolute risk or ARR is aggregating an

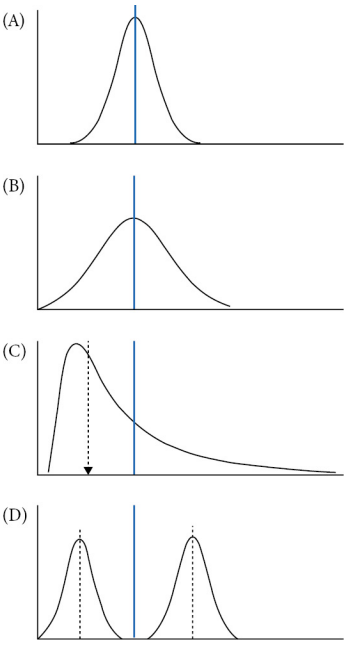



**Figure 6.4** The distribution of individual propensities (or changes in propensity due to an intervention) in four hypothetical populations, represented by four charts graphing the individual propensity (or change in propensity) on the x-axis and the frequency of individuals on the y-axis. The dashed line in C represents the median, while the dashed lines in D represent the mode. All populations have the same mean (represented by the solid lines) but markedly different dispersions.

extreme bimodal distribution (similar to graph D). What kind of distribution typically occurs in medicine?

It is difficult to know because our information about risks comes from aggregate studies in which it is difficult if not impossible to tease apart the influence of causal heterogeneity from the play of chance. We do know (for instance) that the absolute risk of a cardiovascular event will vary widely from less than 1% to over 50% in subgroups of patients (Goff et al. 2014), which is consistent with a large spread of cardiovascular propensities in the population. We have risk calculators that stratify risk, but not down to the level of the individual, and studies of their performance suggest residual variability in individual risk within risk-stratified populations (Sniderman et al. 2015). We also saw in section 6.2 that for some medical interventions,

marked heterogeneity in treatment effects is expected. For instance, cancer therapies will have no effect on risk of death for individuals who die from a mechanism of death not targeted by the therapy, and many interventions, including cardiovascular interventions like endarterectomy and digoxin, have paradoxical harms in addition to their beneficial effects. In general, we should worry that any absolute risk or population risk reduction might be a poor estimate for the individual.[13]

David Kent and Rodney Hayward (2007) argue that the distribution of patient risk in trial populations will often resemble Figure 6.4C (which is adapted from their article). Disease risk factors often cluster; for instance, age is correlated with cumulative exposure to other risk factors, while cardiovascular risk factors are often mutually associated (as with the newly minted diagnosis of cardiovascular-kidney-metabolic syndrome; Ndumele et al. 2023). Thus there will often be a small number of high-risk individuals, creating a large high-risk tail. Meanwhile, the mean risk in the trial population is generally much lower. Because the individual risk cannot be less than zero, the graph cannot be symmetrical around the mean. Thus the graph of patient risk will be right-skewed, and the mean will substantially overestimate risk for most individuals. Kent and colleagues (2016) found that this was indeed the case in a survey of RCTs. If the RRR is simultaneously similar among individuals (the distribution of RRRs resembles graph A), then the distribution of ARRs will also be right-skewed (as in graph C), and the mean will substantially overestimate the absolute treatment effect for most individuals. However, Kent and Hayward (2007) note that if the intervention has paradoxical harms that are not less common among low-risk individuals, then even the relative treatment effect can vary in low-risk compared to high-risk individuals and may be net harmful among those at low risk.

That aggregates might not be a good estimate for the individual is an epistemic problem, a problem with our clinical research methodology or

[13] The original Framingham Heart Study investigators seemed to agree, cautioning that their cardiovascular disease risk estimates were accurate for groups but not necessarily for individuals (AHA 1973), despite Framingham's role in instilling the concept of the patient's individual risk, statistically calculated from group data, in mainstream medicine (Phillips 2025). The developers of a recent risk prediction tool concur, writing, "It is also important to note that all risk estimates are based on population averages and may under- or overestimate risk in any given individual" (Khan et al. 2023, 1990). Meanwhile, clinical trial pioneer Austin Bradford Hill famously remarked that clinical trial results represent "a group reaction. They show that one group fared better than another, that given a certain treatment patients, on the average, get better more frequently or more rapidly, or both. We cannot necessarily, perhaps very rarely, pass from that to stating exactly what effect the treatment will have on a particular patient" (1952, 117).

heterogeneous study populations. However, it is a problem only given an interpretation of epidemiological probabilities that takes them to estimate patient propensities. Therefore, to solve this epistemic problem, we could apply an epistemic fix by studying less heterogenous groups. (This is the approach of so-called precision medicine.) Or we could apply a conceptual fix, an alternative interpretation of epidemiological probabilities that accommodates heterogeneous research studies. I favor this latter strategy because we are still far away from widely realizing the lofty aims of precision medicine (which I discuss in Chapter 9).

Even if a propensity interpretation of the single-case probabilities provided by epidemiological research leads to inaccurate inferences about patient propensities, one might be skeptical that a conceptual misinterpretation of these single-case probabilities could make any practical difference. However, one reason to treat as worrisome the possibility that aggregate risks poorly predict patient propensities is that assuming that the aggregate is a good estimate for individuals has two alarming consequences if we are mistaken. First, we ascribe a non-zero patient risk to everyone (except when the absolute risk is near zero). Because patient risks and risk factors are increasingly treated as diseases in their own right (Schwartz 2008; Greene 2007; Aronowitz 2015), this ascription turns everyone into a patient (e.g., a "cardiovascular risk patient") with an individual risk that might need monitoring and intervention. Moreover, we assign a non-zero patient risk reduction to everyone (except when the effect size is near zero), which makes it appear as though the intervention is universally effective—everyone stands to benefit by having their own risk lowered. The assumption of an equal distribution of patient risk and risk reduction could lead to high healthcare anxiety, unrealistic expectations for treatments, and overtreatment.

This possibility may be more than hypothetical. Aronowitz (2015) points to the increasing use of contralateral prophylactic mastectomy, the surgical removal of a cancer-free breast in patients with cancer in the opposite breast to reduce the risk of future breast cancer. The use of the operation more than doubled between 1998 and 2003. This rise was similar among all stages of non-metastatic breast cancer (Tuttle et al. 2007) and thus relatively independent of both the absolute risk of contralateral breast cancer and the ARR from prophylactic mastectomy measured in epidemiological studies. Aronowitz sees this finding as an example of how people at different levels of risk use decision-making strategies that are relatively insensitive to objective probabilities. It may reflect the tendency to project epidemiological

risks onto individuals and to view individual risk as a target of treatment in a way that does not closely track the magnitude of epidemiological risks. Indeed, in a qualitative study of decision-making among women who had undergone the surgery, Katharine Rendle and colleagues (2015) found that many participants valued having their individual risk lowered, even if this lowering was objectively small.[14] If patients are undergoing treatment because they mistakenly believe that treatment will eliminate their propensity or risk, then the propensity (mis)interpretation of patient probabilities is not benign.

The final problem with the view that epidemiological evidence estimates individual risks-as-propensities—another epistemic problem—is that it is inconsistent with the best probability model for modeling epidemiological probabilities. A 3% patient propensity toward heart attack is a property of the individual. A good probability model for this propensity is a 100-faced die that has a 3% propensity toward rolling light gray because 3 of its 100 faces are light gray. We infer that there is a 3% probability of rolling light gray by studying the die's physical properties (number of faces, weighting, symmetry) or by rolling the die or an identical population of dice many times and making inferences about the probability of rolling light gray for the die from the frequency of light gray rolls. Yet we do *not* infer that a patient has a 3% probability toward a heart attack by studying their physiologic properties. Nor do we infer the patient's probability by studying outcomes in an identical population of patients—patients in the population are similar in certain ways but cannot be assumed to each have an identical propensity. We infer that a patient has a 3% probability because, on average, 3 out of 100 patients who are similar in certain ways will have a heart attack. A 100-faced die is a good probability model for patient propensity, yet inconsistent with how we come to learn about the patient probabilities we infer from epidemiological evidence. This provides good grounds to reject a propensity interpretation of patient-level probabilities inferred from epidemiological evidence.

Recall the argument scheme introduced in Chapter 5 that is commonly used to infer patient probabilities from population risks: Particularization.

[14] Participants were also motivated by the avoidance of future cancer surveillance and worry and by improved or maintained breast appearance (Rendle et al. 2015). In a qualitative study of patients' understanding of colorectal cancer risks, participants usually associated risk with danger and only sometimes with probability (Han et al. 2009), which might motivate patients to neutralize the danger regardless of the probability of harm. Most participants in the study endorsed a frequency-type (or ontic) interpretation of absolute risk, while a minority endorsed a belief-type (or epistemic) interpretation.

One of the assumptions on which this inference relies is what I called the "unbiased population assumption": the probability that a patient will have their prognosis assessed is independent of their risk of the outcome and the probability that they will respond to the intervention in a particular way. Then the probability that a patient will have the outcome is equal to the proportion of individuals in the population who will have the outcome.

Given the unbiased population assumption, we can model the epistemic situation of the treating clinician using as our probability model a lottery in which the clinician draws a ball randomly from an urn containing light gray and dark gray balls. Imagine that the treated and untreated diagrams in Figure 6.1 represent two such urns. When a patient presents to clinic, it is epistemically equivalent to a random draw from one of these two urns. In deciding whether to use a treatment, the patient locates themselves in either the treated urn or the untreated urn. The single-case probability that the drawn ball (patient) will be light gray (have a heart attack) is given by the finite frequency of balls in their urn—in Figure 6.1, the counterfactual probabilities are 3% (untreated) and 2% (treated). We can infer that there is a 3% probability that the ball we have drawn is light gray not from facts about this ball but from facts about the overall urn: the frequency of light gray balls in the urn, along with the assumption that the ball-drawing is not biased toward particular balls. But the probability is still a single-case probability, the probability that a ball that has already been drawn from this urn is light gray. Our evidence for the frequency of light gray balls is provided by counting the colored balls, just as the epidemiologist's evidence for the absolute risk in a study population consists in counting clinical outcomes. The urn probability model is thus consistent with how we come to learn about patient-level probabilities. The next question to ask is *How should we interpret these probabilities, if not as propensities?*

In summary, the patient's risk is most readily interpreted ontically, as a patient propensity. I am not denying that there might be clinical propensities, and the alternative I will now propose is agnostic to their existence. However, there are serious problems with interpreting the patient-level probabilities inferred from epidemiological evidence as propensities: aggregate epidemiological outcomes might be a poor estimate of patient propensities and reductions in patient propensity, misestimation could lead to overdiagnosis and overtreatment of patient risk, and a propensity view is inconsistent with the best probability model for modeling epidemiological probabilities. Rather than patient propensities ("the patient's risk," "the patient's risk

reduction"), medicine should reinterpret patient-level probabilities epistemically.

### 6.3.2 Credence and Epidemiological Uncertainty

Medicine's implicit interpretation of patient probabilities inferred from epidemiological research is ontic: these probabilities represent the patient's risk or the reduction in the patient's risk, where the patient's risk is a propensity, a property of the patient. Assuming that you are a member of the untreated population represented in Figure 6.1, your individual risk or propensity of having a heart attack in the next ten years is 3%. The virtue of a statin is that it lowers your propensity of heart attack from 3% to 2%.

In comparison, an epistemic interpretation of these probabilities sees them as the probability of a hypothesis about the patient. The relevant prognostic hypothesis is *H*: the patient will have a heart attack in the next ten years if untreated. The epidemiological evidence represented in Figure 6.1 can be viewed as providing probabilistic support for this hypothesis: 3% represents the probability of *H* conditional on *E*, where *E* represents the epidemiological evidence plus evidence for the assumptions of the Risk Generalization-Particularization Model discussed in Chapter 5, including evidence allowing us to extrapolate the epidemiological study results to a target population that includes the patient (Generalization) and evidence allowing us to infer a patient-level probability from the study results (Particularization). The propensity view is more aligned with the measurement theory of medical evidence discussed in Chapter 5, in which evidence consists in risk measurements, because propensities are measurable properties. In contrast, an epistemic interpretation is more aligned with the argument theory or inference theory of medical evidence defended in that chapter, in which evidence supports premises in arguments for a medical hypothesis.

There are various epistemic interpretations on the philosopher's menu, including classical, logical, and credence varieties (Gillies 2000; Suarez 2020). For instance, a clinician's or patient's credence in a hypothesis is their coherent degree of belief that the hypothesis is true. A *rational* or objective credence is one that a rational agent would adopt, given (for instance) certain evidence. In the case of epidemiological probabilities, this credence is conditional on learning the epidemiological evidence. I will endorse interpreting patient probabilities inferred from epidemiological evidence as rational

credences over other epistemic interpretations for several reasons. First, credence is currently the leading contender among epistemic interpretations of probability in philosophy, statistics, and medicine. Thinking about credences in prognosis and therapy might also not be a huge conceptual leap for contemporary evidence-based medicine, as Bayesian inference is commonplace in evidence-based diagnosis (Guyatt et al. 2015). Moreover, uncertainty is widely recognized in medicine as an important feature of clinical reasoning and decision-making (e.g., Spiegelhalter 2008; Djulbegovic et al. 2011; Tonelli and Upshur 2019). Credences are a ready and useful way of representing uncertainty. Medical evidence is often seen as having some bearing on medical uncertainty. Epidemiological evidence in particular informs uncertainty by grounding our credences in population frequencies.

More important, a credence interpretation is well suited to our best probability model for epidemiological probabilities: an urn containing colored balls. The clinician has drawn a ball randomly from the untreated urn in Figure 6.1. The probability that they have drawn a light gray ball is 3%, which can be interpreted as a rational agent's credence that the ball is light gray given evidence of a 3% frequency of light gray balls in the urn. I do not mean that this credence is the one that a particular physician will have, descriptively speaking. I mean that it is a rational credence to adopt. Analogously, a 3% probability that a patient will have a heart attack is a rational credence to adopt if the evidence *E* on which Figure 6.1 depends is the clinician's evidence. As I mentioned, *E* includes evidence beyond the clinical trial results. It can be rational for an agent to adopt this epidemiological probability as their credence when they have no additional, independent evidence relevant to *H*. It amounts to accepting that the urn model represents their epistemic situation well.[15]

The greatest benefit of my credence interpretation may be that it avoids the pitfalls of the propensity interpretation discussed in section 6.3.1. There I raised the concern that because their populations are heterogeneous, epidemiological studies might provide an inaccurate estimate of the patient's true individual propensity or the patient's true reduction in individual propensity from treatment. The clinician's rational credence cannot be inaccurate in the same way as the propensity estimate because it is not intended

[15] If we grant these assumptions, the justification for setting their credence to the absolute risk can be provided by one of several rational principles, including the Principal Principle (Lewis 1980), the Principle of Direct Probability (Hacking 1965), or the Relative Principle (Beebee and Papineau 1997). I do not have space to discuss these principles or their application in this context.

as some ideal credence equal to the patient's propensity. It rather represents the clinician's uncertainty given variability in the outcome in the population, captured by the aggregate outcome. Some (or all) of this variability may be due to chance: the medical condition might actually be governed by propensities. Some (or all) of the variability might also result from causal heterogeneity: there could be variability among patients in mechanisms for the medical condition as well as in the medical intervention's ability to disrupt a patient's mechanism given the presence or absence of cofactors.

The study's aggregate outcome—the absolute risk, or the population-level risk reduction—packages together chanciness and causal heterogeneity. Its weakness in estimating a patient's propensity or the reduction in their propensity is simultaneously its strength in serving to represent the physician's uncertainty to the extent that the physician does not know the distribution of individual outcomes in a population (individual risks and individual risk reductions of varying magnitudes) and where in the distribution the particular patient is located. In other words, which of the graphs in Figure 6.4 best represents the distribution of individual outcomes (if any of them do), and where should we plot the patient's data point? Because patient-level credences derived from epidemiological study results generally reflect this characteristic uncertainty, I will give those credences and that uncertainty a name: epidemiological uncertainty.

The second pitfall of the propensity interpretation was that it makes a patient of all of us. Physicians and patients often regard a patient's risk as a medical condition in its own right and seek to eliminate or mitigate it. Attributing an equal propensity or propensity reduction to all members of the target population could lead to overdiagnosis of patient risk and overtreatment when the epidemiological evidence provides a poor estimate for patient propensities. In contrast, there is no risk of overdiagnosing or overtreating patient risk on the credence interpretation because the credence interpretation does not interpret patient-level probabilities as patient risks. It instead interprets patient probabilities as representing epidemiological uncertainty, which has its source in data on the population.

There are also positive advantages of adopting the credence interpretation. For one, it allows us to better conceive of how to combine other kinds of evidence and uncertainties, including diagnostic evidence and uncertainty, with evidence from prognostic and therapeutic studies and with epidemiological uncertainty. Particularization is an inference with empirical assumptions in need of their own evidence. For instance, consider the other

assumption on which *E* relies, the "population membership assumption," which holds that the patient is a member of the target population from which the patient-level probabilities are derived (Chapter 5). If the clinician is certain of this assumption, it might be rational to adopt a credence equal to the absolute risk in the target population. However, if the clinician is *un*certain of this assumption, it might *not* be rational to set their credence to the absolute risk.

To illustrate, according to the ACC/AHA risk calculator used by the statin decision aid in Figure 6.1 (Goff et al. 2014), the untreated absolute risk of a heart attack is roughly twice as great among diabetic patients (Mayo Clinic 2024). Assuming that diabetes does not alter the relative reduction in risk of heart attack with a statin (Costa et al. 2006), then the absolute reduction in risk will also be roughly twice as great among diabetic patients (ARR = RRR x untreated absolute risk). These differences might matter for decision-making. If a clinician is uncertain as to whether a patient has diabetes, they should also be uncertain as to whether it would be better to estimate their patient's risk (and risk reduction) from the absolute risk (and ARR) among diabetic patients or from among non-diabetic patients.

Thankfully, on a credence interpretation they are not forced to choose, and it would be rationally suboptimal to do so. Say that they are 50% certain that their patient has diabetes. Their credence that the patient will have a heart attack *without* a statin should then be equal to the unweighted average of the *untreated* absolute risks (given diabetes, given no diabetes); their credence that the patient will have a heart attack *with* a statin should be equal to the unweighted average of the *treated* absolute risks (given diabetes, given no diabetes).[16] The resulting probabilities don't correspond to any of the propensities estimated by the epidemiological evidence on a propensity

[16] $H$ is the hypothesis that the patient will have a heart attack. $D = 1$ if the patient has diabetes, $D = 0$ if the patient doesn't have diabetes. The clinician's conditional credences $p(H|D = 1)$ and $p(H|D = 0)$ can be set to the untreated absolute risk in the diabetic and nondiabetic population, respectively. $p(D = 1)$ and $p(D = 0)$ are the clinician's credences in the diagnosis of diabetes and in the diagnosis of no diabetes, respectively. Then the clinician's unconditional credence that the patient will have a heart attack can be calculated as follows:

$$p(H) = p(H \mid D = 1) \cdot p(D = 1) + p(H \mid D = 0) \cdot p(D = 0)$$

For example, $$p(H) = p(H \mid D = 1) \cdot 0.5 + p(H \mid D = 0) \cdot 0.5$$

Note that in the case of diabetes, physicians typically have access to diagnostic evidence that can reduce their uncertainty in the diagnosis, thus modifying the values of $p(D = 1)$ and $p(D = 0)$, along with $p(H)$.

interpretation. Therefore a credence interpretation of patient probabilities makes it easier to see how diagnostic evidence and uncertainties can be combined with prognostic and therapeutic evidence and with epidemiological uncertainty.

I mentioned that it can be rational for an agent to adopt epidemiological uncertainty as their credence that a particular patient will have a heart attack when the agent lacks independent evidence (i.e., evidence not used in inferring the epidemiological uncertainty). If they *do* have independent evidence, a credence interpretation provides some general guidance on how the agent can instead proceed: they can treat epidemiological uncertainty as their prior probability, their initial credence, that they can then update using the independent evidence. For instance, the ACC/AHA guidelines on the primary prevention of cardiovascular disease list "risk enhancing factors" (Arnett et al. 2019). These are cardiovascular risk factors that aren't included in the ACC/AHA risk calculator used to predict the absolute risk of cardiovascular events. The ACC/AHA recommends using risk-enhancing factors to adjust risk estimates, particularly for patients at intermediate risk of a cardiovascular event, for whom treatment recommendations are less definitive. It might not always be possible make this correction quantitatively and infer a precise adjusted probability, but even inferring that the patient's probability of a cardiovascular event is *higher* than the calculator predicts could help inform decision-making, particularly when alternative courses of action initially seem equally reasonable.

A propensity interpretation can make sense of this adjustment procedure by viewing it as a matter of correcting our estimate of the patient's propensity. However, risk-enhancing factors, which include a family history of a premature cardiovascular event, are not uncommon. Thus the propensity interpretation would have to concede that epidemiological evidence often provides an inaccurate estimate of the patient's propensity in need of correcting, which undermines the view that we should interpret patient probabilities inferred from epidemiologic evidence as the patient's propensity. In contrast, a credence interpretation views this adjustment as a matter of rationally updating our credence that the patient will have a cardiovascular event (and thus would benefit from a preventive intervention) given evidence that the patient has additional risk factors. A credence interpretation allows us to see how further evidence independent of the epidemiologic study results can inform prognostic and therapeutic predictions.

Arjun Devanasan argues that my credence interpretation "misidentifies the real issue and so fails to offer a meaningful solution to it. . . . [T]he real problem for inferring probabilities in the individual case is the reference class problem" (2023, 1001). The patient belongs to many populations in which the absolute risk may differ. Which absolute risk should the rational agent set their credence to? This is an instance of the infamous reference class problem (Reichenbach 1949). Of course, we lack good data to estimate the absolute risk in many clinical populations. Yet in estimating cardiovascular risk, we sometimes do have good data: the statin decision aid allows the user to choose from several risk calculators that use slightly different risk factors and sometimes differ in the untreated absolute risk they provide for the same patient (Mayo Clinic 2024).

While I have not solved epidemiological medicine's reference class problem here, as Devanasan notes, it is arguably not a unique problem for credence interpretations of single-case probabilities (Hajek 2007), so it does not undermine my interpretation of patient probabilities inferred from epidemiologic evidence. Similarly, because it arguably confronts various interpretations of probability, solving the reference class problem would not help us decide among competing interpretations. Yet decide carefully among interpretations we must if we are to make sense of the relevance of epidemiologic evidence for individual patient probabilities without falling into the epistemic and practical traps that some interpretations—namely, the propensity interpretation—set for us. So while I agree that the reference class problem is a real problem, I disagree that it is the only real problem that dogs single-case probabilities in medicine.[17]

On the other hand, the credence interpretation also brings with it a unique challenge: the relevance of the effect size for clinical decision-making is at first glance more difficult to understand when patient probabilities are

[17] In fact, by not taking seriously the problem of interpreting single-case probabilities in medicine, Devanasan (2023) misdescribes epidemiological medicine's reference class problem by wedding it to the propensity interpretation: "The problem that arises when we try to infer what will happen to an individual based on frequency data in a population is the problem of inferring the propensity for a particular event occurring that involves that individual. When there are multiple reference classes to which that individual belongs, rationality dictates that we take all of them into account when inferring the individual's propensity, and this leads to the reference class problem" (1008). Meanwhile, "[i]f the AR [absolute risk] in a reference class is all we know, then we cannot do any better than inferring the propensity of an individual member of the class to be equal to the AR" (1005). I argue above that we shouldn't infer the patient's propensity from the absolute risk, and the credence interpretation I propose doesn't require us to infer the patient's propensity at all. Correctly characterizing epidemiological medicine's reference class problem and the relevance of the absolute risk for the individual requires taking seriously the interpretation problem.

interpreted epistemically. On the propensity interpretation, the effect size measured in an intervention study estimates a patient's reduction in risk. A reduction in a patient's risk of a harmful outcome is easily seen as a benefit of treatment, the kind of information that is relevant in risk communication and clinical decision-making. An equivalent credence interpretation would understand the effect size as estimating a change in the clinician's or patient's credence in the hypothesis that the patient will have the harmful outcome. But why should they consider using the intervention if all it will do is change their beliefs? Learning further evidence might also change their beliefs—do medical interventions thus play the same role as medical evidence? On an epistemic interpretation of patient probabilities, the causal information captured by effect measures (discussed in section 6.2) doesn't seem to transfer over to the patient probabilities.

While there is no direct equivalent to a patient's risk reduction on the credence interpretation, that does not provide a good reason to endorse the propensity interpretation over the credence interpretation because the patient risk reductions offered by the propensity interpretation of epidemiological evidence might be misleading. However, the puzzle of missing patient-level causal information should force us to dig deeper to unearth how the population-level effect size is relevant for patient-level decision-making on the credence interpretation.

One tantalizing possibility is suggested by the depiction of the jet black person in Figure 6.1. The decision aid's text says that 1 person out of 100 will be "saved from a heart attack" by taking the statin, which is equal to the absolute reduction in risk of heart attack due to a statin in this population (causal ARR = 1%). This might suggest that this one person will have a heart attack prevented by the statin and that, in general, the causal ARR for statins reveals the frequency of individuals for whom the statin prevents a heart attack. We might then infer that the probability that this patient—presenting randomly from the population—will have a heart attack prevented by a statin is 1%.[18] This probability could be interpreted as the clinician's credence in the hypothesis that a statin will prevent a heart attack for their patient. The causal ARR is then relevant for patient-level decision-making about whether to use an intervention because the hypothesis it supports is about patient-level prevention.

[18] In this spirit, Rothwell et al. claim, "An absolute risk reduction tells us what chance an individual has of benefiting from treatment—i.e., an absolute risk reduction of 25% suggests that there is a 1-in-4 chance of benefit" (2005, 257).

However, this reasoning makes assumptions about the causal ARR that might not be true and are not established by the causal ARR alone. The causal ARR simply suggests that a statin will reduce the frequency of heart attacks by 1% on net in this population, not that it prevents a heart attack in exactly one person in the deterministic sense that there will be one individual who would have a heart attack if untreated but would not have a heart attack if treated. That will be true only if the causal variables discussed in section 6.2 as influencing the effect size take on very specific values; for instance, if only 3 people out of 100 have clinically significant atherosclerosis, if the probability of heart attack given the undisrupted atherosclerosis mechanism is 100%, if a statin successfully disrupts atherosclerosis for just 1 of those 3 people, and if it reduces their probability of a heart attack from 100% to 0. Hypothetically, it could instead be, for instance, that there are 6 people with clinically significant atherosclerosis, each with a 50% probability of a heart attack if their atherosclerosis is undisrupted, and that a statin works for 4 of them (the ones with disruption cofactors), reducing their probability from 50% to 25% each. (Perhaps they remain at 25% probability of a heart attack because they had atherosclerotic plaque before taking the statin.) We cannot say for any of these four individuals that they would have had a heart attack without the statin and that it prevented their heart attack, let alone designate exactly one of them as the one who is saved from a heart attack.

As we also saw in section 6.2, sometimes medical interventions have paradoxical harms and help bring about the medical condition that in other cases they help avert. For example, in Figure 6.3, endarterectomy contributes to stroke prevention, reducing the absolute risk of stroke in this population by 5% overall. However, it achieves this overall reduction by increasing the risk of operative stroke shortly after surgery by 7% and reducing the risk of ipsilateral stroke down the road by 12%. Thus it is plainly not correct to say that endarterectomy will prevent a stroke in 5% of patients, nor to adopt a 5% credence in the hypothesis that endarterectomy will prevent a stroke for one of these patients.[19] If we do not wish to rule out the possibility of paradoxical

[19] The effect size might also underestimate the number of individuals in which an intervention disrupted a mechanism due to over-prevention and preventive preemption, in which the disruptive influence of the intervention on the mechanism is made redundant by another mechanism disruptor. For instance, endarterectomy might disrupt the atherosclerosis→stroke mechanism in an individual, but if that individual was already taking a statin, then the statin might also disrupt the mechanism (over-prevention) or might have done so if endarterectomy were not used (preventive preemption). These individuals do not contribute to endarterectomy's effect size for stroke prevention because they will not have a stroke regardless of whether they undergo endarterectomy. Whether these are genuine cases of prevention depends on how "prevention" is defined (Fuller 2022b).

harms or chancy prevention, we must not assume that the causal ARR can generally supply the probability of a hypothesis about patient-level prevention or cure. We might in some cases be justified in ruling out these phenomena using additional evidence or theory. My point is that the causal ARR on its own should not be interpreted this way.[20]

Jan Sprenger and Jacob Stegenga (2017) provide a different way of understanding the relevance of the ARR (and similar "absolute measures") in clinical decision-making that relies on the apparatus of expected utility theory. In short, given the relevant utilities (the magnitude of value and disvalue attached to various outcomes), the ARR allows a decision-maker to determine whether the expected utility (EU) of one intervention (the probability-weighted sum of the utilities) exceeds the EU of another intervention or no intervention. In their terminology, the ARR is "EU-sufficient." (In contrast, the RRR and similar "relative measures" is not EU-sufficient, and one needs to combine it with further information to tell whether an intervention exceeds another in expected utility; namely, one needs to know the absolute risk in the comparison scenario.) Since this determination is the central determination in expected utility theory, allowing a decision-maker to decide which intervention maximizes expected value, we can conclude that on an influential theory of rational decision-making, the ARR is informative for decision-making (and more informative than the RRR).[21] And given some patient-level interpretation of the probabilities represented by the ARR, it is informative for *patient-level* decision-making.

To be relevant for any decision, the probabilities involved in rational decision-making must be counterfactually related. A patient should want to know their probability of the outcome if they receive the intervention compared with an alterative scenario in which they do not receive the

[20] Compared to the patient risk reduction interpretation, the "probability of prevention" (mis)interpretation of the causal ARR may be more benign, especially if it is interpreted as a credence and does not imply that the patient will benefit by having their risk reduced. Moreover, it may be more intuitive to understand compared with alternative ways of communicating the relevance of the causal ARR for clinical decision-making. However, the interpretation may not be completely benign: a clinician or patient might wrongly think that there is a separate probability of harm with respect to the same outcome to weigh against this "probability of prevention" when in fact this "probability of prevention" represents the difference between the intervention's preventive and harmful effects.

[21] Sprenger and Stegenga (2017) provide two further arguments in support of using absolute over relative outcome measures in medicine: relative measures lead physicians and patients to overestimate the effectiveness of medical interventions, and they fail to satisfy several desiderata for a measure of causal strength that absolute measures satisfy. Ina Jäntgen (2023) argues that absolute measures maintain their advantage over relative measures for rational decision-making across a range of information scenarios.

intervention. This comparison is a counterfactual one: if one of the scenarios is actual, the other must be counterfactual because the two scenarios are mutually incompatible. A counterfactual comparison of probabilities is the right kind of comparison for a decision-maker because to make a decision is to choose between two (or more) futures that are counterfactually related to each other. Thankfully, on the interpretation of effect measures I provided in section 6.2, the *causal* ARR measures the difference in absolute risk between two populations that are counterfactually related in the way required for the ARR to be informative for decision-making. And when the ARR measured in an epidemiological study is unbiased, it estimates the causal ARR (Rothman et al. 2008).

To be relevant for *patient-level* decision-making, the counterfactually related probabilities also need a *patient-level* interpretation. On a propensity interpretation, their relevance is straightforward. The choice of which intervention to use is a choice between two counterfactually related scenarios in which the patient's propensity is different. Because the probabilities represent their personal propensity, they are obviously relevant for personal decision-making. How, then, can we understand the relevance of counterfactually related credences when they are decidedly *impersonal*, reflecting the frequency of the outcome in a population?

The answer is that regardless of whether the probabilities are interpreted as propensities or credences, we can model the decision as a choice between two lotteries. When we do so, it becomes apparent that the structure of the decision is the same, even if the probability model is different. Consider once more a causal ARR of 1%. On a propensity interpretation of patient probabilities, the decision-maker faces a choice between two hundred-faced dice that each have some dark gray faces and some light gray faces. The dice are identical except that the "treated die" has one fewer light gray face. If light gray is a harmful outcome, it would be rational (all else being equal) to roll the die with one fewer light gray face. The choice is a counterfactual one: if the decision-maker chose this die, they didn't choose the other die.

On my credence interpretation of patient probabilities, the decision-maker instead confronts the two ball-and-urn lotteries in Figure 6.1. The "treated urn" has one fewer light gray ball. A ball has already been drawn randomly from each urn. (After all, a patient has no choice of who in the population they are.) If light gray is a harmful outcome, it would be rational (all else being equal) to choose the ball from the urn with one fewer light gray ball. The choice is a counterfactual one: if the decision-maker chose this

urn, they didn't choose the other urn. And choosing the ball drawn from the urn with one fewer light gray ball is rationally equivalent to choosing an intervention with a causal ARR of 1%. Of course, if the decision-maker could peek at the identity of the drawn ball, they might be able to better decide based on their knowledge of the ball's identity. (If the outcome is governed by propensities, perhaps different balls are painted with different dyes, and each dye has a certain known probability of turning light gray.) Otherwise, they should make their decision the same way as the decision-maker using the propensity model, by choosing between the two available probabilities associated with the counterfactual lotteries. If it is rational to choose one die over the other in the propensity model, it is rational to choose one ball over the other in the structurally similar credence model.

In summary, we ought to interpret patient-level probabilities inferred from epidemiological evidence as a clinician's or patient's rational credence in a hypothesis about whether the patient will have a clinical outcome. This credence represents their "epidemiological uncertainty" regarding the distribution of individual outcomes in a population and where in the distribution the particular patient is located. The credence interpretation avoids the pitfalls of the propensity interpretation and allows us to see how prognostic and therapeutic evidence and uncertainty can be combined with other kinds of evidence and uncertainty, including diagnostic evidence and independent evidence. In patient-level decision-making, it understands the causal ARR not as estimating an intervention's probability-lowering effect on the individual nor as providing the probability of patient-level prevention or cure but as counterfactually comparing two probabilities in a way that—for decision-making purposes—is structurally similar to the comparison made on the propensity interpretation.

How best to communicate probabilities to decision-makers is a problem for the field of risk communication and is partly an empirical one. Patients often struggle to comprehend risks, risk reductions, and similar quantities (Gigerenzer and Edwards 2003; Schwartz 2011). Less discussed is that they may also struggle with concepts of probability or risk, may interpret probabilities differently depending on the language used to express them, or may fail to understand or appreciate the relevance of population data for themselves as individuals (Han et al. 2009; Weinfurt et al. 2012). As we already saw, clinicians, scientists, and other health experts are also not immune to the difficulties. My main argument in this chapter is a philosophical one, that an epidemiological uncertainty or credence interpretation of patient

probabilities inferred from epidemiological evidence is better theoretically justified than a patient risk or propensity interpretation. Moreover, the theoretical relevance of epidemiological uncertainty for personal decision-making can be modeled using counterfactual ball-and-urn lotteries.

However, my account does suggest tentative proposals for communicating patient probabilities to evidence users and decision-makers in practice when we choose to communicate probabilities. Communicators can describe patient probabilities in epistemic terms, explain how epidemiological evidence informs epidemiological uncertainty, and discuss these probabilities and their relevance to decision-making using analogies to games of chance,[22] especially ball-and-urn lotteries. Meanwhile, they should avoid the language of individual patient risk and patient risk reduction and couching prognosis and the benefit of preventive and curative medical interventions in these terms. The risk communication literature provides suggestions for how decision aids can be used to enhance communication and decision-making, including with icon arrays like the ones in Figures 6.1–6.3 (Zipkin et al. 2014; Stacey et al. 2017; Bonner et al. 2021). The way they are typically constructed, icon arrays best communicate population risks, but they can potentially help communicate epidemiological uncertainty too, especially if they are designed to represent lotteries. As these proposals are philosophically well-grounded, they deserve further empirical investigation.

## 6.4. Conclusion

Preventive and curative medical interventions reduce the risk of medical conditions in a population compared to their counterfactual omission, commonly by disrupting an etiological or pathological mechanism for the condition, but their effects on individuals vary. Thus, rather than inferring some fact about the patient from epidemiological effect measures—the patient's risk or their risk reduction—we should interpret effect measures as informing our uncertainty about the patient's outcome and their response to medical intervention. In contemporary risky medicine, medical interventions lower

[22] David Spiegelhalter (2008) suggests that risks summarize uncertainty, and that uncertainty should be interpreted as a subjective probability (i.e., a credence). He also suggests that making analogies with games of chance could help in communicating uncertainty. In this vein, Jerome Hoffman and colleagues (2006) designed a decision aid in which potential treatment outcomes are represented on roulette wheels. Paul Han (2021) proposes conceptual frameworks for managing medical uncertainty, drawing on philosophical and empirical research.

the absolute risk of a medical condition in a population compared to that same population if untreated. They do so by disrupting mechanisms that cause or sustain the medical condition. Effect measures like the causal ARR aggregate effects on individuals in the population, blurring treatment effect heterogeneity arising from variation in the presence of mechanisms or pathology as well the intervention's cofactors. Patient-level probabilities are often inferred from the absolute risk and from effect measures, and discussed as though they were the patient's propensity to have an outcome or the reduction in their propensity due to the intervention. However, epidemiological evidence might provide a poor estimate for patient propensities, and misestimation could lead to overdiagnosis and overtreatment of patient risk. Thus I proposed to reinterpret patient-level probabilities inferred from epidemiological evidence as rational credences representing epidemiological uncertainty, or uncertainty about the population distribution of individual outcomes and where in the distribution the particular patient is located. This interpretation overcomes the pitfalls of the propensity interpretation, makes clear the relevance of further evidence and uncertainties in forming patient probabilities, and allows us to model the causal relevance of the population effect size for patient-level decision-making using counterfactual ball-and-urn lotteries.

In the next chapter, I turn to the problems of extrapolating the effect size from epidemiological evidence and evaluating the external validity of clinical trials.

# 7
# Problems of Extrapolation


## Summary

Assessing the external validity of a clinical trial requires solving the problem of extrapolation, the question of how we can extrapolate the trial's causal findings to a target population. I argue that the problem of extrapolation actually consists of *two* problems: a logical problem of extrapolation and an evidential problem of extrapolation. While the logical problem is concerned with finding good extrapolation arguments, the evidential problem is concerned with finding good evidence for those arguments. I scrutinize existing solutions to these problems and develop my own solution to the logical problem based on my analysis of preventive and curative medical interventions, as well as my own solution to the evidential problem that utilizes structural analogies.

The logical problem of extrapolation involves constructing an extrapolation argument, which will contain premises asserting causal similarity between the study and target populations. Meanwhile, the evidential problem of extrapolation involves identifying the kinds of evidence that would in principle confirm those premises. Even a partial solution to these problems must overcome the extrapolator's circle, the challenge of finding premises and evidence that would allow us to extrapolate but that would not allow us to infer the causal conclusion about the target population without the study results, thus rendering moot the need to extrapolate. I examine solutions to the problem of extrapolation proposed by others, arguing that they are generally solutions to the logical problem only, even my own solution that relies on an analysis of how a preventive or curative medical intervention works, which I call mechanism-cofactor analysis. I then reject evidence-based medicine's solution to the evidential problem, simple extrapolation, in which we extrapolate the relative treatment effect in the absence of a compelling reason not to extrapolate. I propose a new solution to the evidential problem in which structural similarities between factors in the study and the target are used to establish the relevant causal similarities. None of these solutions



*The New Modern Medicine*. Jonathan Fuller, Oxford University Press. © Oxford University Press 2025.
DOI: 10.1093/9780190066178.003.0008

addresses the practical problem of extrapolation: when to accept the conclusion of our extrapolation argument. Though I don't resolve that problem here, any plausible solution cannot demand utmost certainty in our extrapolation inference.

## 7.1. The External Validity of Clinical Trials

Carotid endarterectomy is the surgical operation in which atherosclerotic plaque is removed from the carotid artery to relieve carotid stenosis, or narrowing in the artery, and (hopefully) prevent a future stroke. In 1991, European and North American research groups tested the effectiveness of endarterectomy for stroke prevention in randomized trials (European Carotid Surgery Trialists' Collaborative Group 1991; H. Barnett et al. 1991). They compared endarterectomy to no treatment in patients with symptomatic stenosis—that is, patients with carotid stenosis who'd had a recent stroke or similar ischemic event—and advanced carotid narrowing of 70–99%. The results were positive. In the North American trial, the relative reduction in risk of ipsilateral ischemic stroke—stroke on the same side as the artery operated on—was over 60% (H. Barnett et al. 1991). By extrapolating these results, clinicians might then predict a relative risk reduction (RRR) of over 60% in future patient populations. Yet doing so would involve a leap, from the rigorous comfort of the trial population to the uncertain external world. Without establishing external validity, such an inference would be a leap of faith.

As it turns out, it would land us in trouble. It was later discovered that the effectiveness of endarterectomy varies widely. A meta-analysis of clinical trials found that the relative risk reduction was less impressive for patients with 50–69% carotid stenosis (RRR = 25%; Rothwell et al. 2003). Another analysis of people with greater than 50% stenosis found that endarterectomy was even less effective for women (RRR = 16%) and for patients undergoing endarterectomy more than twelve weeks after their most recent stroke event (RRR = 15%), among other groups (Rothwell et al. 2004). For women treated more than twelve weeks after their most recent stroke event, the treatment was actually net harmful (a relative risk *increase* of about 350%; Rothwell et al. 2005). This example is not isolated. In Chapter 5, I introduced other examples of non-generalizability involving HIV antivirals and COVID-19 vaccines. Of course, the fact that a few trials were not generalizable to

important target populations does not tell us much as philosophers or physicians trying to extrapolate from countless other trials. What we really want to know is how, in general, an evidence user can predict when trial results can be extrapolated. This question is a specific instance of "the problem of extrapolation" (Steel 2008): How can we extrapolate causal conclusions made in the context of a study to a particular target setting?[1]

In Chapter 5, I analyzed the Risk Generalization-Particularization Model, which is the standard model of therapeutic prediction in the new modern medicine of today. Extrapolation or Generalization is the first stage in the model: predict whether the therapeutic effect size measured in the study would also obtain in the target setting. As I discussed, this inference relies on a sufficient similarity assumption. In their book *Evidence-Based Policy*, Nancy Cartwright and Jeremy Hardie (2012, 46) argue that the advice that the target population should be "sufficiently similar" to the study population is unhelpful ("you are lucky if it gets you anywhere"). One of their worries is that this advice is vague. As I suggested, the sufficient similarity assumption is often unarticulated. My analysis of treatment effect heterogeneity associated with preventive and curative medical interventions from Chapter 6 will help in articulating the assumption. However, though analyzing the sufficient similarity assumption is necessary, it is not enough to solve the problems of extrapolation that I will explore in this chapter. That is because it provides a solution to only one problem of extrapolation, what I will call "the logical problem of extrapolation." While most existing solutions address the logical problem, none successfully answers what I will call "the evidential problem of extrapolation."

"External validity" is a term used in clinical epidemiology and evidence-based medicine (EBM) to describe the extent to which the study results (e.g., the effect size measured in the study) accurately predict the results of intervening in some target setting. For decades, medical commentators have worried that the external validity of many clinical trials for important target populations is poor. As Richard Kravitz and colleagues once put it, "Although randomized controlled trials (RCTs) usually are good experiments, they

[1] I will mostly be concerned with the problem of extrapolating the effect of a given intervention on a given outcome from a study to a target that might differ in terms of characteristics of the population and setting, though in 7.4.2 I will also touch on the possibility that the intervention itself might differ. I won't consider cases where the outcome differs (e.g., the study outcome is a surrogate for the target outcome) or where the relevant interventions were not compared head to head in the same study, which are additional ways in which the study might provide only "indirect" evidence for the effect in the target (Guyatt, Oxman, Kunz, et al. 2011).

often are poor surveys" (2004, 667). Meanwhile, Peter Rothwell once claimed that "[l]ack of consideration of external validity is the most frequent criticism by clinicians of RCTs" (2005a, 82).[2] Because trials have often excluded typical patients through narrow eligibility criteria (Van Spall et al. 2007; Kennedy-Martin et al. 2015; Tan et al. 2022) and because they have often been run in specialized settings in which the intervention and ancillary care might be atypical (Rothwell 2005a), they may not be representative of real-world target settings. The external validity of clinical trials is an important problem for contemporary evidence-based, epidemiological medicine because of its heavy reliance on clinical trial evidence and because of these concerns.[3] Yet Judea Pearl, for one, suggests that the present scientific culture is "paralyzed" by threats to external validity and "looks with suspicion and disbelief on the very idea that threats can be disarmed" (Pearl and Mackenzie 2018, 357).

We will see that there are several obstacles that any theoretically informed solution to the problem of external validity must disarm, and that it generally cannot succeed without representing mechanisms and causes. One such obstacle is known as "the extrapolator's circle," which Daniel Steel (2008) argues any adequate mechanisms solution to the problem of extrapolation must overcome. The extrapolator's circle is the challenge of specifying what information would allow us to extrapolate from the study without that information already containing the conclusion we hoped to extrapolate. Not everyone agrees that this problem can be successfully resolved by a "mechanisms approach" to the problem of extrapolation (Howick et al. 2013), nor that a solution to the problem of extrapolation needs to resolve it (Marcellesi 2015). To settle these disagreements, we must first start by disambiguating the problems of extrapolation.

The problem of extrapolation is actually two problems of extrapolation: a logical problem and an evidential problem. The logical problem demands a good extrapolation argument, while the evidential problem demands to know what evidence could support those premises in the argument warranting study-target comparability. I develop my own solution to

[2] More recently, Kennedy-Martin et al. write, "That the external validity of RCT results is often limited is widely acknowledged by clinicians as a problem when it comes to extrapolating data to the patients seen in everyday practice" (2015, 10).

[3] Cartwright and Hardie (2012) as well as Julian Reiss (2019) argue that thinking of the task of predicting effectiveness as one of assessing the external validity of evidence is a poor framing of the problem. I agree that we should adopt a broader and more flexible picture of medical evidence, like that offered by Cartwright and Hardie's argument theory of evidence, which I endorsed in Chapter 5. But because external validity occupies so much medical attention, it is worth examining in this chapter.

the logical problem, mechanism-cofactor analysis, based on my analysis of preventive and curative medical interventions, as well as my own solution to the evidential problem that utilizes structural analogies. In section 7.2, I distinguish the two problems of extrapolation and show how they are both threatened by the extrapolator's circle. In section 7.3, I scrutinize solutions to the logical problem that rely on our knowledge of causes and mechanisms. In section 7.4, I reject a solution to the evidential problem known as "simple extrapolation" in favor of a new proposal utilizing structural analogies. In section 7.5, I dispel the worry that these solutions require a high level of certainty in our knowledge of the relevant mechanisms by introducing a third problem of extrapolation, the practical problem.

## 7.2. Two Problems of Extrapolation

Steel (2008) uses the term "the problem of extrapolation" to refer to the challenge of determining whether causal relationships that hold in a base or study population can be extrapolated to a target setting. It is a problem whenever the effects of an intervention are heterogeneous among individuals because then the effect size and even the direction of the net effect will depend on which individuals comprise the population. In this section, I will explain why *the* problem of extrapolation should be viewed as *two* problems of extrapolation: a logical problem and an evidential problem.

Alexandre Marcellesi (2015) calls Steel's problem the "problem of external validity." While internal validity describes the accuracy of measurements of the study population, "external validity" refers to the extent to which the study results accurately predict the results in some target setting. Unlike internal validity, external validity is not a property of the study itself that we can judge by examining the study in isolation because it depends on certain similarities obtaining between the study and the target. Also unlike internal validity, external validity might vary with the target population. So it is best to think of external validity as a relationship that holds between the study and a particular target. Subgroup analysis is commonly discussed as a solution to the problem of external validity (Rothwell 2005b). In subgroup analysis, the treatment effect is measured in a subgroup of the study. We might then extrapolate the subgroup-specific effect to the target population when the target population shares characteristics with the subgroup. While external validity may be more likely to obtain between a subgroup and the target

population, we must still determine whether this is indeed the case, which is the problem of external validity or problem of extrapolation articulated by Marcellesi and Steel. Throughout this chapter, I will discuss the problem of extrapolation without distinguishing whether we are extrapolating from a subgroup or the overall study population.

The science of extrapolation[4] is less well developed compared to the science of causal inference, in which causal conclusions are inferred about the study population from associations measured in the study. Nonetheless, Marcellesi (2015) argues that the problem of external validity is no longer an unsolved problem. According to Marcellesi, "the problem of external validity has been solved" (1317) by Cartwright and Hardie (2012) and independently by Pearl and Bareinboim (2014).[5] Cartwright and Hardie's account, for example, succeeds because it "gives a clear and unambiguous answer to the question of the circumstances in which two populations are 'similar' enough that one can extrapolate a causal claim from one to the other" (Marcellesi 2015, 1314). For Marcellesi, a solution to the problem of external validity lists sufficient conditions or circumstances under which the causal effect in the target will be similar to the causal effect in the study. For instance, it provides a general formulation of the sufficient similarity assumption discussed in Chapter 5, which is invoked in "generalizing" a treatment's effect size. These general sufficient circumstances can be made less abstract, given enough understanding of how a specific intervention produces a specific effect.

One ingredient that Cartwright and Hardie list in their recipe for predicting whether an intervention will work is "support factors," causes that interact with the intervention to bring about the effect. The distribution of support factors matters when extrapolating because they determine whether the intervention will work for an individual, and thus their distribution in the population determines the magnitude of effect in the population. A sufficient similarity assumption should specify in general terms how the distributions of support factors in a study and target must compare for the effect size to "travel" from one population to the other.

[4] For examples of philosophical work on the problem of extrapolation in the health sciences, see LaFollette and Shanks (1995), Cartwright (2007, 2010, 2011), Steel (2008), Howick et al. (2013), Clarke et al. (2014), Stegenga (2015), Baetu (2016), Lemoine (2017a), Reiss (2019), Eersel et al. (2019), Wilde and Parkkinen (2019), La Caze and Osimani (2020), and Fuller (2021b).

[5] Marcellesi (2015) is referring to what he calls "predictive external validity inferences," in which effects are prospectively predicted in a target population, rather than "explanatory external validity inferences," in which effects in a target population are explained. I too will be focusing on predictive external validity inferences.

Marcellesi suggests that Cartwright and Hardie's account may fail to break the extrapolator's circle, a daunting task that Steel sets for any hopeful solution to the problem of extrapolation. The extrapolator's circle is the difficulty of having the right kind of causal information that would allow us to extrapolate from the study but would not allow us to predict the effect in the target more directly, making the extrapolation redundant. I will further discuss the extrapolator's circle momentarily. Marcellesi (2015, 1315) argues that a good account of external validity like Cartwright and Hardie's need not bother with this problem: "[T]his account gives the correct answer to the question of the circumstances under which predictive EV [external validity] inferences are good. Assuming that this claim is true, what more could one want of an account of EV inferences?"

However, there is an additional problem in predicting the effect in the target population beyond laying out an external validity inference or argument, namely, how do we judge the soundness of the argument? What evidence can we use to support the premises on which the inference relies? In Cartwright and Hardie's solution, how do we know whether the intervention's support factors are distributed similarly in both populations, especially when our understanding of how the intervention works is incomplete? Analogously, general accounts of causal inference often state the assumptions that are sufficient for inferring causation from association (e.g., the intervention is associated with the outcome even when all confounders are controlled), but it is a separate question how to establish those conditions (e.g., through a randomized trial versus by stratifying data on known confounders). As in the causal inference analogy, the details of the sufficient similarity assumption will determine what counts as evidence for the assumption, but there may be more than one potential evidential strategy, and it might not be immediately clear or uncontroversial what strategies are available and which are most promising.

Marcellesi is aware of this additional challenge; he distinguishes between "analyses" of "the conditions under which EV inferences are good inferences" and "methods for drawing such inferences" (2015, 1315).[6] He even concedes that these "methods" might be confronted with the extrapolator's circle. But he argues that all we should expect from a solution to the problem of external validity is an "analysis," and analyses need not show how to break the circle.

[6] Likewise, Steel distinguishes between the "basis of an analogy and types of evidence that could be given to show that the basis is present" (2010, 1060).

To determine whether Marcellesi's claim is reasonable, it is necessary to further distinguish the two solutions that Marcellesi describes, which map onto two distinct problems.

Instead of *the* problem of extrapolation, it is more helpful to speak of *two* problems of extrapolation, which I will call "the logical problem of extrapolation" and the "evidential problem of extrapolation." The logical problem of extrapolation is the problem of articulating a good extrapolation argument. Because attempts to articulate an extrapolation argument in philosophy and science are often formal or semi-formal, the arguments offered in solution to the logical problem are often deductive and contain sufficient conditions for their conclusion, including a premise establishing that the study and target are similar in certain ways. Because the effect size is a causal property, this premise will usually refer to causal circumstances.

By stipulating that the extrapolation argument must be a "good argument," I don't mean that it must be sound—that would require all our good extrapolation arguments to be infallible. But a good extrapolation argument must at the very least be non-fallacious, avoiding formal and informal logical fallacies. When the argument is meant to be deductive (as are all the extrapolation arguments that will concern us here), it must avoid formal logical fallacies so that the argument is deductively valid. Shortly I will argue that a good extrapolation argument or serial argument should also not contain any dispensable premises, including those premises describing the study results. Put differently, the argument should not remain valid in the absence of those premises.

In contrast, the evidential problem of extrapolation is the problem of determining what data would in principle constitute good evidence or support for those premises in an extrapolation argument that establish the similarity of the study and the target. This support will often take the form of evidence collected directly from the study and target, but in section 7.4.1 I will provide an example of an approach called "simple extrapolation" that instead relies on more general assumptions about the transportability of effect sizes. One immediately obvious way to provide support for those premises claiming that the populations are comparable is to use representative sampling when implementing the study, for instance, by sampling randomly from the target population. Then the extrapolator would require limited knowledge of the relevant causal circumstances beyond information required to judge that their statistical inference is sound. However, random sampling of trial populations is rarely (if ever) done. Thus I will concentrate on more

theoretical solutions to the evidential problem, those that harness knowledge of the relevant causes as well as data on their distributions in the two populations.[7]

By stipulating that the relevant evidence must be "good evidence," I don't mean that the evidence must be definitive or even strong. It can provide only some support for the premises or render them merely probable. But it must be evidence for those premises in the extrapolation argument that establish comparability, *and* it must not support premises in another argument that warrants the conclusion of the extrapolation argument more directly by ignoring the study results. It will become clearer what I mean in a moment.

Recall the inference theory or argument theory of evidence introduced in Chapter 5. To warrant a medical hypothesis, the inference theory requires us to first lay out an argument for the hypothesis. Data or facts are then evidence when they establish a premise in that argument (or a premise in a sub-argument). These two components of the inference theory—arguments and data—map onto the two problems of extrapolation, the logical and the evidential. The problems of extrapolation are thus a specific application of the inference theory to extrapolation inferences.

Extrapolation arguments can be formulated in different ways with different premises and conclusions, as we will see in the next section. But they all follow a general pattern. All extrapolation arguments that involve extrapolating a causal effect include the following kinds of premise and conclusion:

A. Premise about the effect in the study.
B. Premise(s) about the comparability of the study and target.
C. Conclusion about the effect in the target.

[7] There have been enthusiastic calls for pragmatic clinical trials in medicine (e.g., Zwarenstein and Treweek 2009). Pragmatic trials are often designed to be representative of typical target populations, not by randomly sampling but by deliberately enrolling typical patients in typical care settings. They often have permissive eligibility criteria for participation and are integrated into routine care. The trial design is basically meant to support the premises in our extrapolation argument warranting study-target comparability. Can pragmatic trials then avoid the need for theoretical knowledge? The answer is no (or not entirely) because designing the trial to be representative of typical patients and typical care settings requires some information about the causal variables that would make the trial *relevantly* representative. The trial will not be similar to the target population in every conceivable way, and some kinds of similarity would be unimportant anyway. What matters is that the trial is similar to the target in the right—causally relevant—ways. So even in extrapolating from pragmatic trials we cannot escape the need to consider pertinent causal similarities and differences. Pragmatic trials relocate much of this consideration to the trial design stage rather than the trial analysis stage. See Tresker (2022) on problems with conceptualizing pragmatic trials and with the "pragmatic trial"/"explanatory trial" distinction.

The logical problem of extrapolation involves inserting the relevant premises and conclusion into this scheme, while the evidential problem involves determining what evidence in principle would support the premises. While there might be several desiderata for a good solution to the problems of extrapolation, I argue that even a good *partial* solution to each problem must overcome the extrapolator's circle.

The extrapolator's circle is a problem explored extensively by Steel.[8] To establish causal similarity between two populations, one could compare the distribution of the causes in the study to their distribution in the target. But if we knew what the relevant causes of the outcome were and how they were distributed in the target, then we might already be able to predict the effect in the target without extrapolating from the study. Steel: "The challenge posed by the extrapolator's circle, then, is to explain how it is possible to establish the similarity of the model [study] and target without already knowing what one wants to extrapolate" (2010, 1059). We can say that the extrapolator's circle is "violated" (or use similar language) when an extrapolator fails to successfully answer its challenge. The challenge is less daunting for solutions that warrant comparability in ways that require minimal causal knowledge such as those that use random sampling. However, for solutions that require more detailed causal understanding, the circle looms larger.

We can understand the extrapolator's circle better through the distinction between the two problems of extrapolation I just described. Just as there are two problems of extrapolation, there are two extrapolator's circles, or at least two versions of the extrapolator's circle. The logical version of the extrapolator's circle is violated whenever premises describing the study results are dispensable. In the general argument scheme above, if B includes a premise or premises that allow us to infer C while ignoring A because those premise(s) and C together make a valid argument, then we can infer the conclusion without extrapolating from the study results referenced in A, and they are dispensable.

For instance, a premise $B_S$ might reference the identity of the relevant causal factors and their distribution in the *study* as well as the relationships among the causes and the effect of interest in the *study*. A premise $B_T$ might

[8] The problem is also discussed by Hugh LaFollette and Niall Shanks (1995) in their work on animal extrapolation, in which they argue that it is insurmountable in this context; it is one of the reasons that we cannot justifiably extrapolate from animal models to humans. As a reductio ad absurdum against their argument, Steel (2008) notes that because the structure of the problem is the same when extrapolating from human studies to human target populations, their argument implies that we can never extrapolate from human studies either—an unacceptable implication.

reference the identity of the relevant causal factors and their distribution in the *target* as well as the relationships among the causes and the effect in the *target*. Comparing premises $B_S$ and $B_T$ side by side might be enough to establish that the two populations are causally similar in ways that guarantee a similar effect. But then we might also be able to infer the effect in the target—conclusion C—from $B_T$ alone, in which case we violate the extrapolator's circle because conclusion C along with a subset of the premises not including premise A together make a valid argument or serial argument. The argument is not circular in the informal logician's sense because C is not included among the premises. It is "circular" in the looser sense that B contains all the information that one needs to infer C, and thus if we knew B we would already know C.

Unfortunately, the extrapolator's circle will not always present itself to an extrapolator in such an obvious way, which requires us to consider a second, "evidential" version of the extrapolator's circle. Say that instead of $B_S$ and $B_T$, B includes a single premise B* claiming that the distributions of the relevant causes and their relationships with the effect—whatever those causes and relationships might be—are similar in the study and the target. Because B* does not include information about the identity of the causes and how they bring about the effect, it is plausible that we cannot infer C from B* alone. However, recall that the evidential problem of extrapolation requires us to determine what evidence in principle would support B*. Let's say that the evidence supports B* indirectly by supporting another argument in which $B_S$ and $B_T$ are the premises and B* is the conclusion. Recall that we can infer C from $B_T$ directly. Thus, if we have evidence for all the premises in the argument for B* (including $B_T$), we also have evidence for all the premises in a third argument, in which $B_T$ is the sole premise and C is the conclusion. We can say that the evidential version of the extrapolator's circle is violated here in the sense that our (indirect) evidence for B in the extrapolation argument supports C more directly in another argument that does not include A.

In the two versions of the extrapolator's circle, the fundamental problem is the same: we don't need premise A or evidence for premise A to infer conclusion C because the evidence and other premises that would allow us to infer C from A also allow us to infer C *without A*. We don't need to extrapolate from the study, and we can ignore the study results. In the evidential extrapolator's circle, the subset comprised of B* and C (and not A) does not form a valid argument, which was the way that I defined the logical extrapolator's circle. However, because B* is warranted in another argument by $B_S$ and $B_T$, we

could easily replace B* with $B_S$ and $B_T$ in the extrapolation argument, and then we *do* violate the logical extrapolator's circle. Because we can generally alternate between logical and evidential formulations of the extrapolator's circle, it is best to view them as two versions of the same fatal flaw.

Having accepted roughly the distinction I draw between the logical problem of extrapolation and the evidential problem of extrapolation, Marcellesi (2015) suggests that (1) the extrapolator's circle confronts only the evidential problem; (2) all we should expect from a solution to the problem of extrapolation is a solution to the logical problem; therefore (3) we should not expect a solution to the problem of extrapolation to grapple with the extrapolator's circle. However, we should not accept any of these three suggestions.

I've shown that a logical version of the extrapolator's circle *does* confront the logical problem of extrapolation. As we have just seen, and as I will further demonstrate through concrete proposals in the next section, solutions to the logical problem of extrapolation can often avoid the logical version of the extrapolator's circle. But these concrete proposals do not include a solution to the evidential problem of extrapolation, which raises the worry that their solution avoids the logical version of the circle only by pushing the challenge posed by the extrapolator's circle onto the evidential problem of extrapolation, in the way that the generic scheme above avoided the logical extrapolator's circle by substituting $B_S$ and $B_T$ with B*, only to have the extrapolator's circle come back to bite when it was time to provide evidence for the premises. Thus, even if a solution to the logical problem of extrapolation evades the extrapolator's circle, we should also demand a solution to the evidential problem to ensure that the difficulty of the extrapolator's circle has not simply been postponed.

Given the seriousness of the extrapolator's circle, we should demand that any good solution to the problems of extrapolation show how it can be avoided rather than postponed (if indeed it can be avoided). The problem of extrapolation is the problem of how we can extrapolate from a study in those cases when we cannot measure or predict the causal effect in the target more directly (without extrapolating). A solution that allowed us to predict the effect in the target population without extrapolating would not be a solution to the problem of extrapolation; it would be an alternative to extrapolation. Thus in the rest of this chapter I will seek to provide a partial solution to both problems of extrapolation. Though my proposal is partial, it shows how a paired solution to both the logical and evidential problems can successfully escape the extrapolator's circle.

The problems of extrapolation are not singular; they are two: a logical problem and an evidential problem. The logical problem asks us to formulate a good extrapolation argument, while the evidential problem asks us to determine in principle what evidence could support its premises. Solving the problems of extrapolation requires overcoming the extrapolator's circle, a flaw in which the premise reporting the study results in the extrapolation argument is dispensable or in which the evidence for the premise(s) establishing study-target comparability could be used to warrant the conclusion in another argument without any need for the study results. In the next section, I start with the logical problem.

## 7.3. The Logical Problem of Extrapolation

The logical problem of extrapolation challenges us to formulate a good extrapolation argument, including premises that warrant taking the study and the target population to be comparable. A good extrapolation argument must avoid the logical version of the extrapolator's circle: the premise describing the study results must not be dispensable for a valid argument. There are different conclusions one could extrapolate from a study, corresponding to different extrapolation arguments. Over three subsections, I sequentially examine three proposals, finding that they provide more or less adequate solutions to the problem of formulating good predictive extrapolation arguments. These solutions are comparative process tracing (Steel 2008, 2010), transport formulas (Pearl and Bareinboim 2014), and mechanism-cofactor analysis.

### 7.3.1 Comparative Process Tracing

In one way or another, theoretically driven solutions to the problems of extrapolation often rely on an understanding of relevant mechanisms (e.g., Clarke et al. 2014). For instance, the mechanistic strategy is used by the International Agency for Research on Cancer (2012) in assessing potential human carcinogens (Cogliano et al. 2008; Wilde and Parkkinen 2019) and was advanced by Steel (2008) with these examples in mind, which typically require extrapolating from animal studies in which an animal model is exposed to the potential carcinogen. However, we can also use a mechanisms

approach to extrapolate the results of a clinical trial in which a human population is exposed to a medical intervention.

Steel's (2008) mechanisms approach to extrapolation is called "comparative process tracing." In comparative process tracing, the extrapolator compares the study with the target to identify relevant mechanistic similarities and differences. The central example Steel uses to illustrate the approach involves a carcinogen named aflatoxin B1 that causes liver cancer in some species but not in others owing to differences in the efficiency of a liver enzyme responsible for its metabolism. In that example, the question is whether we can extrapolate the exposure's carcinogenic effect from certain animals to humans. In studying extrapolations from clinical trials, I will discuss another example that Steel uses elsewhere, the example of efavirenz (Park et al. 2023). I introduced this example in Chapter 5. Efavirenz is an antiretroviral agent for the treatment of chronic HIV infection that was shown in clinical trials done in the U.S. to lower HIV viral load (in combination with other antiretrovirals) with rare neuropsychiatric side effects. However, when used in Zimbabwe, it produced common and significant neuropsychiatric symptoms, including hallucinations and mood changes (L. Nordling 2017). Patients experiencing these side effects stopped taking the drug, which led to ineffective control of their viral load. Thus not only did the effect size for significant neuropsychiatric symptoms in the trials fail to generalize to the Zimbabwean population, but so did the effect size for management of viral load. What went wrong?

Like most drugs, efavirenz is metabolized in the liver as part of its elimination from the body. One of the main liver enzymes responsible for its metabolism is CYP2B6, which lowers efavirenz's blood concentration. Mutations in the gene for CYP2B6 can lead to a defective enzyme that fails to efficiently metabolize efavirenz, causing efavirenz's blood levels to rise until they are high enough to cause neuropsychiatric symptoms in the brain. Park and colleagues (2023) represent the metabolic pathway using a causal directed graph. I have produced my own, modified graph (Figure 7.1).

Causal directed graphs can be used to model the connectivity of a mechanism, or the qualitative causal relationships among its variables. For instance, in Figure 7.1, the variable *I*, representing efavirenz ingestion, causes variable *B*, representing certain blood levels of the drug, which in turn causes variable *S*, representing neuropsychiatric symptoms. The arrows represent

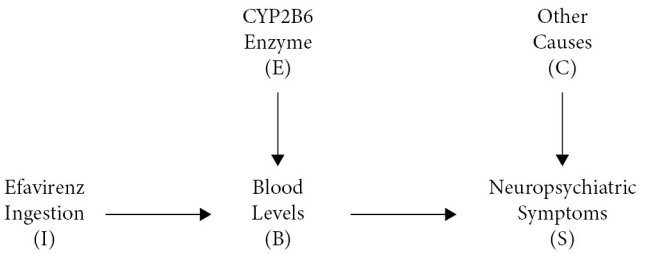



**Figure 7.1** Mechanism of efavirenz's neuropsychiatric side effects, represented using a causal directed graph.

causal dependence and point from cause to effect.[9] A "directed path" or "causal path" is a chain through a series of variables that can be traced by following the direction of the arrows, such as $I \rightarrow B \rightarrow S$ in Figure 7.1. A "mediator" of two variables is located somewhere on a causal path between them. In Figure 7.1, B is a mediator between *I* and *S*. For a more detailed overview of the use of causal diagrams in the health sciences, the reader is encouraged to consult Greenland and colleagues (1999).

For simplicity, we can assume that each variable in Figure 7.1 takes on two values, 0 or 1 (0 for the absence of the cause in a given individual, 1 for its presence in a given individual). Also simplifying matters, we can assume that all the arrows in Figure 7.1 are positive, or that when each variable increases from 0 to 1, it promotes or positively influences (rather than inhibits) its downstream effects. When efavirenz is ingested ($I = 1$) and the mutant CYP2B6 variant is present ($E = 1$), it promotes high blood levels of efavirenz ($B = 1$), which in turn promotes symptoms ($S = 1$). We can pair a causal graph with a probability distribution over its variables, as well as a set of functions relating the probability of each variable to other variables. I will introduce this added element in section 7.3.2, which allows us to derive transport formulas for predicting the effect size.

Steel uses comparative process tracing to extrapolate the "causal relevance" of an exposure. An exposure is causally relevant when it makes a positive (or

[9] For ease of illustration, I have omitted an arrow that points from side effects back to efavirenz ingestion. That arrow would represent the influence of those side effects in causing patients to stop ingesting efavirenz.

negative) causal difference with respect to the outcome, for instance, raising (or lowering) the outcome's probability.[10] Steel's Extrapolation Theorem connects claims about causal relevance to mechanisms represented in causal graphs. According to the Extrapolation Theorem, an exposure variable is causally relevant for some outcome variable in a population if and only if there is an undisrupted mechanism, an instantiated causal path, from exposure to outcome for at least some members of that population.[11] In Figure 7.1, efavirenz ingestion is positively causally relevant for neuropsychiatric symptoms when $I = 1$, $B = 1$, and $S = 1$.

The tight connection between causal relevance and causal paths is the basis for comparative process tracing. Say that investigators run a clinical trial of efavirenz and find that some patients experience side effects; in other words, efavirenz is positively causally relevant for side effects in some trial participants. They find that those patients who experience side effects have elevated blood levels of efavirenz. A clinician later considers prescribing efavirenz. But, learning of these results, they worry that efavirenz might cause these side effects in their patient and want to predict whether that will be the case. According to comparative process tracing, efavirenz will be positively relevant for those side effects if their patient shares the same causal path as those unfortunate trial participants. Thus, to extrapolate positive causal relevance, the clinician can conduct a node-for-node comparison to see if each stage in the causal path from ingestion of efavirenz to neuropsychiatric symptoms is undisrupted in their target patient as it was in the trial subgroup experiencing the side effects.

However, if the clinician knew that the entire causal path was undisrupted for their patient, they would already know that efavirenz ingestion has positive causal relevance for the side effects without having to extrapolate from the study because they could infer positive causal relevance from the presence of the undisrupted causal path by the Extrapolation Theorem. The (logical) extrapolator's circle would be violated because they could infer the conclusion of positive causal relevance directly from a premise claiming that the efavirenz→symptoms path is undisrupted. Contrariwise, if we assume

[10] Probability raising is a special case of Steel's definition of positive causal relevance that applies to dichotomous variables. It is the only case I will discuss. In Steel's definition of positive causal relevance, the value of the exposure variable is set by an "ideal intervention" in James Woodward's (2003) sense.

[11] According to Steel, the Theorem assumes "consonance," that there are no other mechanisms from exposure to outcome that have an opposite effect on the outcome.

that the causal path in Figure 7.1 is the only path through which efavirenz causes those symptoms, then upon finding that the mechanism is *disrupted* or inoperative in their patient, the clinician could conclude that efavirenz does not have positive causal relevance for their patient—again, without extrapolating from the clinical trial.

To avoid the extrapolator's circle, Steel proposes a strategy for selective mechanistic comparison that does not require checking every stage of the mechanism in the target but rather those stages that are most likely to be disrupted: "First, learn the mechanism in the [study population], by means of process tracing or other experimental means. . . . Second, compare stages of the mechanism in the [study population] with that of the target [population] in which the two are most likely to differ significantly" (2008, 89). In addition to other desiderata, one should check the farthest causally downstream mediator that is likely to differ because when a causally downstream mediator is present, then its upstream causes must also be present and don't need to be checked (Steel 2008). In the efavirenz pathway, mediator *B* plausibly satisfies this criterion.

Upon finding similarity in critical mechanism stages, and assuming that there are no relevant mechanistic differences downstream of where the extrapolator has checked, the extrapolator can extrapolate the causal relevance of the exposure from the study to the target. We can represent comparative process tracing as an extrapolation argument, using the efavirenz mechanism as our example:

1. STUDY MECHANISM: In the study, $I \rightarrow B \rightarrow S$.
2. PARTIAL TARGET MECHANISM: In the target, $I \rightarrow B$.
3. DOWNSTREAM SIMILARITY: There are no relevant downstream mechanistic differences between study and target.
4. REMAINING TARGET MECHANISM: Therefore, in the target, $B \rightarrow S$ (from 1–3).
5. FULL TARGET MECHANISM: Therefore, in the target, $I \rightarrow B \rightarrow S$ (from 2 and 4).

Premises 1 and 2 are established empirically by studying the mechanism directly in both settings. DOWNSTREAM SIMILARITY is also an empirical assumption. By "relevant downstream mechanistic differences," it means that the mechanism downstream of *B* is either undisrupted in both settings or it is disrupted in both settings. Combining Steel's Extrapolation Theorem

with STUDY MECHANISM, we get the claim that *I* has positive causal relevance for *S*. We get the same claim by combining the Theorem with FULL TARGET MECHANISM. Thus, by combining the Extrapolation Theorem with Steel's comparative process tracing, we can extrapolate positive causal relevance from study to target.

Comparative process tracing is a solution to the logical problem of extrapolation because it lays out the conditions for a valid extrapolation argument. It avoids the logical extrapolator's circle because without the premise describing the study results, the (serial) argument is invalid. Premise 1, describing the trial results, is needed to validly infer premise 4, and premise 4 is needed to validly infer the conclusion. DOWNSTREAM SIMILARITY is not strong enough to warrant the conclusion directly because it doesn't include specific detail about downstream mechanism nodes in the target. However, this comes at a cost: it is not immediately clear what evidence would support DOWNSTREAM SIMILARITY and whether that evidence could avoid the *evidential* extrapolator's circle. Steel (2008) provides no general solution to the evidential problem and the evidential circle.

There is a further problem with comparative process tracing worth discussing, a problem with the method as a solution to the logical problem of extrapolation. Namely, comparative process tracing is limited in its relevance for predictive extrapolation inferences. In a predictive extrapolation inference, the extrapolator hopes to predict whether the intervention will have a certain effect before using the intervention. This prediction can then inform their decision about whether to implement or recommend the intervention. Unfortunately, comparative process tracing is able predict whether the intervention will cause the effect only *after* the intervention is used because it requires checking for a mediator downstream of the intervention. In order to tell whether efavirenz ingestion has resulted in high blood levels, efavirenz must have already been ingested. Comparative process tracing cannot be used to extrapolate in advance of using the intervention.[12]

[12] That does not mean that comparative process tracing is useless for predicting. It may sometimes be possible to discontinue or adjust the intervention after checking for the presence of a mediator before the intervention causes downstream harm or before it is fully implemented. In the efavirenz example, it may not be too late to avert the neuropsychiatric side effects after detecting high blood levels of the drug, for instance, by discontinuing the drug or using an antipsychotic medication. For a cholesterol-lowering drug meant to prevent a cardiovascular event, an important mediator is a lowered blood cholesterol (Borén et al. 2020). After prescribing a statin, a physician can often check that the LDL-cholesterol is lower a few weeks later, and, if it isn't, they can adjust the drug or the dose or counsel the patient regarding use.

Thankfully, a modest modification to comparative process tracing would make it a more suitable approach to predictive extrapolation inferences. This modification suggests itself when we examine efavirenz's metabolic mechanism beyond the causal path from drug to symptoms. The reason why efavirenz ingestion results in high blood levels of the drug in some individuals is that those individuals have the deficient variant of CYP2B6. This variant is not caused by efavirenz; its presence or absence is independent of efavirenz ingestion. Thus the extrapolator can check for the CYP2B6 variant as a marker for future high blood concentration were efavirenz to be used, and they can do so before recommending or implementing efavirenz.[13] Generalizing from this example, rather than a mediator between the exposure and outcome, the extrapolator can check for the presence of nodes off the causal path from exposure to outcome that causally interact with the exposure to bring about the mediator. Steel (2010) discusses a similar proposal in which one checks for a marker of a mediator, such as an effect of that mediator. However, to be maximally useful for predictive extrapolation arguments, that marker should be a cause of the mediator rather than an effect because an effect of a mediator can be detected only after the mediator occurs and, thus, after the exposure occurs. In analyzing the efavirenz example, Steel makes use of another solution to the problem of extrapolation that *does* revolve around the CYP2B6 variant (Park et al. 2023), which I will describe in section 7.3.2.

By suggesting a way that the premises of comparative process tracing's extrapolation argument can be revised, this modification still leaves us with a solution to the *logical* problem of extrapolation only. Moreover, there remains an important limitation of comparative process tracing, and that is the modesty of the result it is designed to extrapolate. Positive causal relevance is a very modest claim, implying in the efavirenz example only that efavirenz raises the probability of its side effects in certain patients with the CYP2B6 variant. This claim does not tell us how much efavirenz raises the probability, whether it does so by a little or a lot, or how often it does so in a population. However, the magnitude of an intervention's effect is one key determinant of whether that intervention is worth using. Where possible, we want a more precise estimate of the effect size. Thus I will examine two further solutions

[13] This application of the strategy I'm proposing is an instance of pharmacogenetics, the tailoring of drug prescribing to an individual using genetic information that allows us to predict how they will respond to the drug. As in the efavirenz example, often this information is about their genotype for a metabolic enzyme.

to the problem of extrapolation that allow us to make a prediction about the magnitude of effect.

### 7.3.2 Transport Formulas

The second solution to the problem of extrapolation I will canvass was developed by Judea Pearl and Elias Bareinboim in a series of articles. Pearl and Bareinboim (2014) derive transport formulas that—given the right circumstances and data—allow the extrapolator to predict the causal effect in the target from probabilities measured elsewhere. Their approach is similar to comparative process tracing in relying on causal graphs. However, unlike comparative process tracing, their transport formulas allow their user to make predictions about the intervention's quantitative effect.

Strictly speaking, Pearl and Bareinboim's approach is not a solution to "the problem of extrapolation" or "external validity" as the problem is commonly framed because their approach is not concerned only with asking whether the same effect holds in two settings. They use the term "transportability" to refer to the computability of causal effects in the target setting from effects measured in a different setting. Their approach requires first asking whether it is possible to compute the overall causal effect in the target population from effects measured in the study. To answer that question, we draw a causal diagram representing the mechanism in both populations. If the diagram reveals that transportability is possible, we can then determine what probabilistic information is needed from both populations and how it can be combined to compute the effect in the target population. To calculate the effect in the target, Pearl and Bareinboim derive transport formulas in the *do*-calculus (Pearl 2009).

I will illustrate the use of transport formulas using our ongoing example of efavirenz. As in every application of Pearl and Bareinboim's strategy, we first draw a causal diagram like the one in Figure 7.1 representing the study and target populations, taking note of where relevant differences between settings occur. Relevant differences between settings include differences in the probability of causal variables and differences in the structural equations relating those causal variables to one another. In the case of efavirenz, the main relevant difference was of the former kind: the probabilities of the two variants of the CYP2B6 enzyme that metabolizes efavirenz differed in the Zimbabwean population compared to the population in

the clinical trials.[14] It turns out that transportability is possible in principle given this kind of difference. The relevant transport formula is the following: $p^*(S=1|do(I=1)) = \sum_E p(S=1\,|\,do(I=1),E) \times p^*(E)$ (Pearl and Bareinboim 2014; Park et al. 2023). In this formula, probabilities with an asterisk are from the target population, while probabilities without an asterisk are from the study population. The transport formula tells us how to calculate the "causal effect" in the target, or the probability of efavirenz's neuropsychiatric side effects ($S = 1$), given an "intervention" that ensures all members of the target take efavirenz ($do(I = 1)$). It calculates this causal effect in the target from the causal effects in the *study* stratified by enzyme variant $E$, multiplied by the probability of $E$ in the *target*. We can represent this approach using the following extrapolation argument:

1. *The causal graph in Figure 7.1 and structural equations governing it are the same in the study and target population, without exception.*
2. *The probability of each causal variable in Figure 7.1 is the same in the study and target population, except for the probability of E (this excludes mediators).*
3. Therefore we can use the following transport formula:
   $p^*(S=1|do(I=1)) = \sum_E p(S=1\,|\,do(I=1),E) \times p^*(E)$
   (from 1, 2, and Theorem 2 in Pearl and Bareinboim 2014).
4. $p(S=1|do(I=1),E=1)=a$, $p(S=1|do(I=1),E=0)=b$ (from the study); $p^*(E=1)=c$, $p^*(E=0)=d$ (from the target).
5. Therefore $p^*(S=1|do(I=1)) = a\times c + b\times d$ (from 3 and 4).

The first two (italicized) premises establish study-target comparability.[15] Like my modification to comparative process tracing in section 7.3.1, they home in on the CYP2B6 enzyme, which causally interacts with efavirenz to

[14] We can indicate this difference by drawing an arrow from a "selection variable" $V$ to $E$ in Figure 7.1, where $V$ allows us to switch from the study ($V = 0$) to the target ($V = 1$), changing the probability of $E$ while holding all else constant. Pearl and Bareinboim (2014) call the resulting causal diagram a "selection diagram."

[15] If there are no relevant causal differences, the causal effect in the overall study population is "directly transportable" to (equal to the causal effect in) the overall target population (Pearl and Bareinboim 2014). In clinical medicine, we ultimately want to predict effects for a target individual. That individual will be located either in the target subpopulation stratified by the wildtype enzyme or in the subpopulation stratified by the mutant enzyme. Rather than the overall target population, it is reasonable to make inferences about the effect in whichever subpopulation the individual occupies. Thus, rather than use the transport formula above, the extrapolator should directly transport the effect in the relevant study subpopulation to the relevant target subpopulation. However, if they do not know the appropriate subpopulation in which to place the individual (perhaps because a test that can

influence its effect on neuropsychiatric symptoms. The nature of the causal differences in the example allows us to use the transport formula in premise 3 to infer the causal effect in premise 5. To do so, we input the probabilities *a*, *b*, *c*, and *d*, if known. The "causal effect" $p^*(S=1|do(I=1))$ is the probability of symptoms in the target population when treated with efavirenz. In order to calculate the effect size, we must compare this probability with the probability of symptoms when not treated with efavirenz, $p^*(S=1|do(I=0))$. To calculate the untreated probability, we can run the argument again, this time measuring the stratified causal effect in the untreated study group.

Pearl and Bareinboim's approach is a solution to the logical problem of extrapolation because it allows us to construct a valid (serial) extrapolation argument. The argument avoids the logical extrapolator's circle because the premise about the study results, premise 4, is indispensable if we wish to infer the conclusion. Pearl and Bareinboim (2014) do not provide a solution to the evidential problem of extrapolation because they don't tell us what evidence would warrant the first two premises about study-target comparability. In particular, they do not solve the evidential extrapolator's circle: how is it possible to establish using evidence that the structural equations are the same and the probability distributions for the causal variables are distributed as required without that evidence being strong enough to calculate the effect in the target directly from the structural equations and probabilities?

A strength of Pearl and Bareinboim's transport formulas is that they can sometimes allow us to predict the effect in the target rather than simply determine whether the effect is the same as in the study, in contrast with the Generalization argument introduced in the next subsection, which can do only the latter. On the other hand, because their solution is domain agnostic, its guidance for those wishing to extrapolate the effects of medical interventions is rather general. In section 7.3.3, I propose a different solution to the problem of extrapolation that is more tailored to medicine in two ways. First, it is based partly on an analysis of how preventive and curative medical interventions work. And second, it explores the conditions under which treatment effects will be "generalizable" to a target population when they are quantified using measures that are commonplace in medicine.

distinguish enzyme subtypes is not available to them), they could then locate the individual in the wider target population and use the transport formula indicated.

### 7.3.3 Mechanism-Cofactor Analysis

In addition to Pearl and Bareinboim (2014), Marcellesi (2015) argues that Cartwright and Hardie (2012) have also solved the problem of extrapolating the effect size. In this subsection, I draw on Cartwright and Hardie's solution along with my analysis of preventive and curative medical interventions from Chapter 6 to propose my own solution to the logical problem of extrapolation, which I call "mechanism-cofactor analysis." Mechanism-cofactor analysis can tell us when a medical intervention's effect size will have external validity and when it will not.

According to Cartwright and Hardie, for the effect of an intervention to be transportable from one setting to another, the same "causal principles" must be true of both settings. A causal principle describes the causal relationship between a particular cause and a particular effect in an individual. For instance, what other factors does the cause need to produce its effect? Causal principles can be qualitative, like the causal graph in Figure 7.1, which shows that efavirenz causes neuropsychiatric symptoms but not how probable those symptoms are given efavirenz. Causal principles can also be quantitative, in which case the probability or value of the outcome is a function of the cause, as with the structural equations in section 7.3.2. Causal principles are often described with some amount of abstraction, omitting some causes or under-specifying others. Figure 7.1 omits further causes of high blood levels of efavirenz, such as overdose, while the variable "other causes," representing further causes of neuropsychiatric symptoms, is nonspecific. These factors are essentially left out of the causal principle, treated as background causes. Consequently, a target individual could have the wild-type CYP2B6 enzyme and yet, unlike similar patients in the efavirenz trials, they could still develop symptoms because the qualitative causal principle represented in Figure 7.1 is not true of them. (Perhaps they have overdosed on efavirenz.)

Now say that we wanted to predict the probability of symptoms for an individual and there were multiple CYP2B6 variants in a population (P. Wang et al. 2019), each conferring a different probability of symptoms. A target individual could have a CYP2B6 variant and yet their probability of developing symptoms could differ from patients with a different CYP2B6 variant because the relevant quantitative causal principles, the structural equations, might not distinguish among variants, and while those principles might accurately predict the probability for a different patient with a different variant, they might not be accurate for the target patient.

The other ingredient in Cartwright and Hardie's account of transportability is "support factors" or "helping factors," so named because they help the intervention produce its effect. An intervention's support factors causally interact with the intervention the way that CYP2B6 interacts with efavirenz. The idea that causes interact with one another to cause an effect is a central component of epidemiology's sufficient-component cause theory (Rothman 1976), introduced in Chapter 3 as a multifactorial theory of the etiology of health outcomes. In SCC theory, causes like efavirenz and CYP2B6 are called "component causes" to distinguish them from "sufficient causes." A sufficient cause consists of multiple interacting component causes. A sufficient cause is sufficient for causing its outcome (causing it with a probability of 1 assuming determinism, causing it with a probability of less than 1 assuming indeterminism). Meanwhile, a component cause is insufficient for causing its outcome. There are often alternative sufficient causes of the same health outcome, for instance, the catch-all variable "other causes (*C*)" in Figure 7.1, which represents component causes that interact within other mechanisms to bring about hallucinations and mood changes (component causes such as substances and mental disorders). When discussing an index component cause, we can say that the component causes with which it causally interacts are its "cofactors" (Greenland and Brumback 2002).[16] For example, the CYP2B6 variant is a cofactor for efavirenz ingestion.

We can represent sufficient causes using causal pies, as in Figure 7.2, showing the efavirenz mechanism. I have simplified sufficient cause I by leaving out myriad other cofactors for efavirenz, including other factors involved in its absorption, distribution, metabolism, and elimination (ADME) in the body.

In Figure 7.2, the pies represent sufficient causes, while the pie slices represent component causes. Causal pies provide an alternative to causal graphs for representing mechanisms. Like causal graphs, causal pies are qualitative causal principles, telling us in non-quantitative terms what causes what. While causal graphs represent causal mediation, causal pies represent causal interaction. Causal graphs include component causes, but conventionally they do not show which causes causally interact with each other (though see C. R. Weinberg 2007 on representing interaction in causal graphs). On the

[16] Component causes of the outcome that are part of sufficient causes that do not include an intervention are potential "confounders" of the association between the intervention and the outcome measured in an epidemiologic study (Fuller 2019).

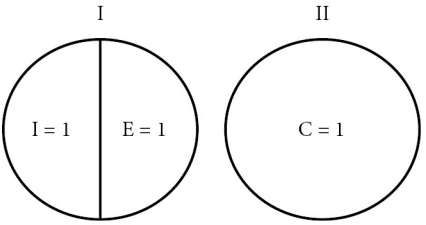



**Figure 7.2** Sufficient causes of neuropsychiatric symptoms. I = 1: efavirenz ingestion, E = 1: presence of CYP2B6 variant, C = 1: presence of other causes or mechanisms.

other hand, causal pies do not show mediators of the component causes included in a sufficient cause.[17]

Our running example has been one in which a medical intervention, efavirenz, causes side effects (which in turn reduces its effectiveness by reducing patient adherence to the intervention). To come up with an extrapolation argument scheme that would also allow us to extrapolate the effect size for a medical intervention's intended effects, I will draw on my account of a preventive or curative medical intervention from Chapter 6. Preventive and curative medical interventions reduce the risk of a medical condition in a population compared to their counterfactual omission the way that carotid endarterectomy reduces the risk of stroke in RCTs. They commonly do so by disrupting a mechanism for the condition or by removing some structural pathology. Endarterectomy does both: it removes atherosclerotic plaque from carotid arteries, which disrupts the etiological mechanism by which this plaque can cause stroke. (The mechanism typically involves a vulnerable plaque causing the formation of a blood clot that travels to the brain, clogs an intracranial artery, and deprives the brain of blood; Mechtouff et al. 2021). A disease mechanism could theoretically produce its outcome with a probability of 100% in the individual, or it could theoretically be chancy, producing its outcome with a probability of less than 100%. Thus, by disrupting the mechanism in the individual, an intervention could theoretically prevent or cure the condition, thwarting an outcome that would otherwise have occurred, or it could merely reduce the risk of the outcome.

[17] Greenland and Brumback (2002) and VanderWeele and Robins (2009) discuss connections between causal graphs and SCC models. See also Cartwright et al. (2021) for a similar joint framework.

If the individual has a second mechanism for the condition that remains undisrupted after intervention, they may even remain at risk.

A preventive or curative medical intervention will generally not work in everyone. For starters, it will not work in those who lack a mechanism targeted by the intervention. Some individuals may even have a mechanism for the same condition, but one that the intervention does not target. For instance, endarterectomy does not target and thus will never disrupt the mechanism by which atrial fibrillation causes ischemic stroke. (The mechanism involves incomplete contraction of the atria of the heart, leading to blood stasis and formation of a blood clot that is then circulated to the brain.) Even in those with a targeted mechanism, the intervention will only disrupt the mechanism given "disruption cofactors" that causally interact with the intervention. (For instance, myriad factors determine the success of a surgery.) In some individuals, the intervention may even produce a paradoxical harm, causing the outcome that in other individuals it prevents. These unlucky individuals have "harm cofactors" that causally interact with the intervention. For some people, endarterectomy causes operative stroke, stroke occurring within the first thirty postoperative days (Rothwell et al. 2003). Harm cofactors responsible for operative stroke or death may include risk factors such as plaque ulceration as well as certain comorbidities (Halm et al. 2009). The presence or absence of preventive and curative medical interventions, their cofactors (disruption cofactors and harm cofactors), and disease mechanisms (those targeted by the intervention as well as those not targeted) are all component causes within sufficient causes of the medical condition.

The effect size for a preventive or curative medical intervention's intended effect is the aggregate of its effects on all individuals in the population. Therefore it is determined by the population's distribution of sufficient causes that contain these component causes (or their absence). Two populations will share a similar effect size whenever they share a similar distribution of disease mechanisms and intervention cofactors. Otherwise they may not share the same effect size.[18]

[18] I say that they *may* not share the same effect size, for two reasons. First, as I will discuss, there are scenarios in which some of these component causes can vary and yet the effect size will be similar. Second, an intervention could theoretically produce the same effect size in two populations but through completely different distributions of causes. This latter possibility will not concern us here because I am interested in how we can extrapolate from a study using our understanding of the causes responsible for the effect in the study.

We can now represent the Generalization argument introduced Chapter 5 as follows, with its sufficient similarity assumption filled in and italicized:

In the study population, the effect size for outcome *Y* due to intervention *X* is *r*.

*X's cofactors and Y's mechanisms are distributed similarly in target population F and the study.*

*Similar causal principles governing X and Y are true of both the study and target population F.*

In target population *F*, the effect size for outcome *Y* due to intervention *X* is *r*.

This extrapolation argument can be used to infer the effect size for an intervention's intended preventive/curative effects, or the effect size for its side effects. For example, to extrapolate the effect size for neuropsychiatric side effects (*Y*) attributable to efavirenz (*X*) from a clinical trial, the distribution of the CYP2B6 variant—efavirenz's cofactor—and of other mechanisms producing neuropsychiatric symptoms must be similar. To extrapolate the effect size for stroke (*Y*) attributed to endarterectomy (*X*), the distribution of endarterectomy's disruption and harm cofactors, as well as mechanisms causing stroke, must be similar. In both cases, similar causal principles must also obtain in both settings, meaning *quantitative* causal principles.[19]

We can call this solution to the logical problem of extrapolation *mechanism-cofactor analysis*. To construct a good extrapolation argument that would allow us to extrapolate the effect size, we first identify the relevant component causes that were responsible for the intervention's net effect in the study population. In extrapolating the intervention's preventive or curative effects, we identify those mechanisms responsible for the disease outcome as well as the intervention's disruption factors and harm cofactors. We then compare their distribution in the study to their distribution in the target. If their distribution is similar, the effect size will be similar. If the distribution of one or more of these component causes differs, then the effect size might also differ, depending on the way we are measuring the effect size and other features of the scenario.

[19] On a counterfactual SCC theory (Suzuki et al. 2018), the quantitative causal principles are the potential outcomes equations relating *Y* to its component causes (see the appendix).

Two commonly reported measures of the effect size are the causal absolute risk reduction (ARR) and the causal relative risk reduction (RRR). Both are derived from the absolute risk, the proportion of individuals in the population with the outcome. The causal ARR can be defined as the difference in absolute risk when comparing the population under intervention to the population under non-intervention (or under an alternate intervention) in a counterfactual comparison. The causal RRR can be defined as the ratio of the causal ARR to the absolute risk under non-intervention (or alternate intervention; Rothman et al. 2008):

causal ARR = absolute risk (under non-intervention) – absolute risk (under intervention)
causal RRR = causal ARR / absolute risk (under non-intervention)

The causal ARR and causal RRR are both determined by the population's distribution of disease mechanisms and intervention cofactors, though because the causal ARR and causal RRR are calculated differently, these causes influence the causal ARR and causal RRR in different ways. (See the appendix for a formal demonstration.) When intervention cofactors or disease mechanisms differ in their strength or severity or in their frequency in the target population compared to the study population, one or both of the causal ARR and causal RRR might also differ.

In general, differences in the intervention's cofactors will produce differences in both the causal ARR *and* the causal RRR. (Both will be greater when harm cofactors are less common and when disruption cofactors are more common among those with a targeted mechanism.)[20] The intervention's cofactors generally influence the causal ARR because the causal ARR captures the difference that the intervention makes in the population, and its cofactors enable it to add or subtract outcomes from the population under intervention. They also generally influence the causal RRR because they influence the causal ARR without influencing the absolute risk under non-intervention. The efavirenz-CYP2B6 example was one in which the frequency of an intervention's harm cofactor differed in the target population. Metabolic enzymes like CYP2B6 and other ADME

[20] In this discussion, I generally assume that differences in other causal variables do not accompany the changes described. If they do, then these accompanying differences can counteract the trends discussed.

factors frequently act as harm or disruption cofactors for drugs by activating them, inactivating them, absorbing them into the body, distributing them to their active sites in the body, or eliminating them from the body (Lai et al. 2022). For all interventions, patient adherence to the intervention is another disruption cofactor, while concurrent interventions frequently act as disruption or harm cofactors through intervention-intervention interactions.

Turning to disease mechanisms, in general, differences in mechanisms not targeted by the intervention will lead to differences in the causal RRR but *not* the causal ARR. (The causal RRR will be less when non-targeted mechanisms are more common.) Non-targeted mechanisms generally don't influence the causal ARR (unless they substitute for a targeted mechanism) because they are not disrupted by the intervention and so are found in equal numbers in the population under intervention and under non-intervention. However, they generally influence the causal RRR because, while they don't generally influence the causal ARR, they do generally influence the absolute risk under non-intervention. Non-targeted mechanisms are a concern for non-specific outcomes like death, stroke, or "flu-like illness" that can be caused by various mechanisms, not all of which are targeted by the intervention. For instance, according to the PREDICT calculator, among women treated with surgery for non-metastatic breast cancer, the proportion dying from causes other than breast cancer tends to increase more with age than does the proportion dying from breast cancer. Accordingly, the causal ARR for death with trastuzumab, an adjuvant cancer treatment, is predicted to vary little (if at all) among older versus younger women, while the causal RRR is predicted to vary greatly (National Health Service 2024).

Meanwhile, with exceptions, differences in mechanisms targeted by the intervention will result in differences in both the causal ARR and the causal RRR. (Both will be greater when targeted mechanisms are more common.) Targeted mechanisms generally influence the causal ARR because they are disrupted by the intervention (assuming the intervention works) and so are found in fewer numbers in the population under intervention. By influencing the ARR, they also generally influence the causal RRR. For example, both the ARR and the RRR for stroke with endarterectomy are greater among patients with 70–99% carotid artery stenosis compared to patients with 50–69% stenosis, a difference driven largely by the greater frequency of carotid plaque-generated thromboembolic strokes among those

with greater stenosis, which endarterectomy disrupts (Rothwell et al. 2003, 2004). In Figure 7.3, this is reflected in the greater frequency of ipsilateral stroke without endarterectomy in those with 70–99% stenosis, represented by individuals who "Benefit."

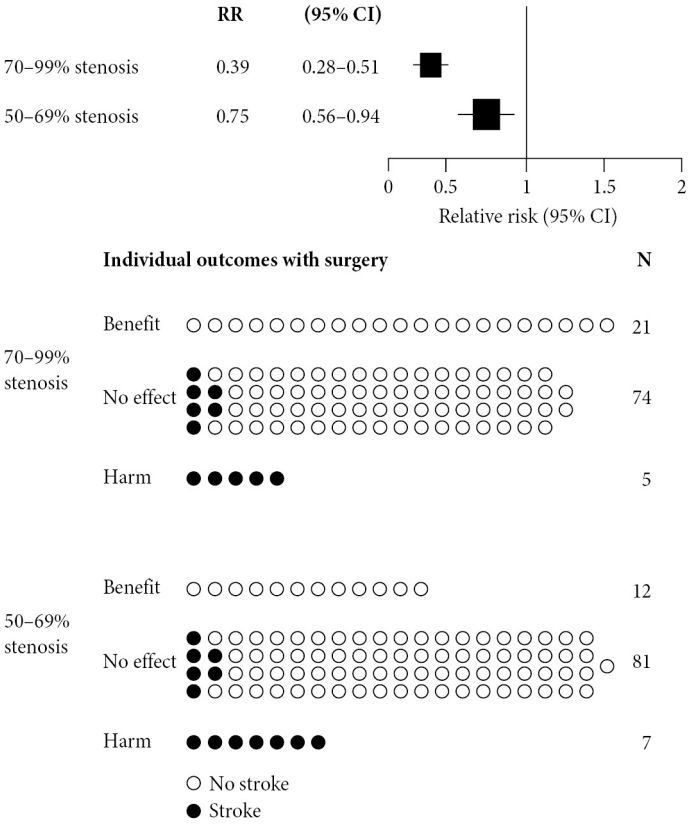



**Figure 7.3** Effectiveness of carotid endarterectomy for five-year stroke prevention among those with different degrees of symptomatic carotid stenosis. The top panel is a forest plot representing the pooled relative risks from a meta-analysis of RCTs (Rothwell et al. 2003). The bottom two panels represent the effect of endarterectomy on individuals. Individuals who "Benefit" would have only an ipsilateral stroke without endarterectomy; individuals who experience "Harm" would succumb to operative stroke or death only with endarterectomy; individuals who experience "No effect" would have the same outcome (ipsilateral stroke, no ipsilateral stroke) whether treated or untreated.
Rothwell et al. (2005, 257).

There is an important exception to this general rule. The causal RRR will *not* be influenced by differences in targeted mechanisms in cases where non-targeted mechanisms and paradoxical harms are rare while the probability of the outcome is similar among those with a targeted mechanism.[21] In these cases, the causal ARR will still increase when targeted mechanisms are more common. But assuming no differences in disruption cofactors, so that the intervention disrupts a similar fraction of all mechanisms, the causal RRR will be unchanged. This might explain why in clinical trials the ARR for heart attacks and strokes with cholesterol-lowering statin drugs increases with an increasing untreated absolute risk, while the RRR is relatively unchanging (Chou et al. 2022).

These principles are helpful in evaluating the common guidance that, when possible, one should calculate the ARR in a target population from the RRR measured in therapeutic research and the untreated absolute risk measured in prognostic research (e.g., using a risk calculator; Glasziou and Irwig 1995; Straus et al. 2011; Post et al. 2013). The motivation for this guidance seems to be that the ARR will be more sensitive to differences in the untreated absolute risk between two populations compared to the RRR (Guyatt, Oxman, Kunz, et al. 2011), and trial populations frequently have a different (often, lower) untreated absolute risk compared to typical target populations (Kennedy-Martin et al. 2015). Prescribing statins for cardiovascular risk reduction is perhaps the example where this guidance is most often applied (Dorresteijn et al. 2011) and serves as a powerful exemplar for its general application. It underlies prediction models like the Statin Choice Decision Aid (Mayo Clinic 2024) as well as clinical guidelines that use cardiovascular risk thresholds for recommending statins (Grundy et al. 2019; Mangione et al. 2022).

Stepping back from the statin example, in general, differences in the absolute risk could be due to differences in targeted mechanisms, non-targeted mechanisms, or both. Mechanism-cofactor analysis reveals that, depending on which mechanisms are varying as well as what other causes are in play, it could be that increases in the absolute risk will increase, decrease, or leave unchanged the causal ARR; the same goes for the causal RRR. Therefore this guidance should consider the impact of disease mechanisms and

[21] If the probability of the outcome varies among those with a targeted mechanism, then disruption cofactors must be equally common among those with each type of targeted mechanism (see the appendix).

intervention cofactors on the generalizability of effect measures across different levels of risk. In principle, a causal RRR could be as sensitive or more sensitive to differences in the untreated absolute risk.[22]

The Generalization argument represented above is similar to Cartwright and Hardie's (2012) "effectiveness argument," translated into the language of cofactors and mechanisms. As an extrapolation argument, it is a solution to the logical problem of extrapolation. It manages to avoid the logical extrapolator's circle because no subset of the argument that excludes the first premise describing the study results is a valid argument. The sufficient similarity assumption does not reveal the identity of the component causes nor the causal principles that would allow us to calculate the effect size from their distribution in the target. However, this leaves it a mystery how an extrapolator could provide evidence for the assumption without knowing in detail the identity of the causes and principles and inviting the extrapolator's circle. While Cartwright and Hardie (2012) provide suggestions for how to support these assumptions with evidence, they do not discuss the extrapolator's circle. Therefore their proposal is not an adequate solution to the evidential problem of extrapolation. In section 7.4.2, I provide a solution to the evidential problem that uses structural analogies to evade the extrapolator's circle.

In summary, the three proposals examined in section 7.3 are all solutions to the logical problem of extrapolation. They all utilize knowledge of causes and mechanisms, along with formal tools and theory. And they each allow us to extrapolate different causal information from a study. Together, they provide a useful toolbox for extrapolating. But none of them gives an adequate solution to the evidential problem of extrapolation because none of them solves the evidential extrapolator's circle. In the next section, I examine solutions to this second problem of extrapolation.

[22] In an analysis of eighteen intervention-outcome pairs studied in clinical trials, the relative treatment effect showed no consistent trend across levels of untreated absolute risk (and showed statistically significant variation for only one pair), while the absolute treatment effect displayed an increasing trend with increases in the untreated absolute risk (Kent et al. 2016). However, there were individual intervention-outcome pairs for which the relative effect showed an increasing trend or a decreasing trend with increases in the untreated absolute risk, and pairs for which the absolute effect showed no trend. A different proposal than the one I've been considering is to stratify trial populations into subgroups defined by the untreated absolute risk (Kent and Hayward 2007), as was done in this analysis.

## 7.4. The Evidential Problem of Extrapolation

The evidential problem of extrapolation requires an extrapolator to determine what would constitute good evidence for an extrapolation argument. Even a partial solution to this problem must overcome the evidential version of the extrapolator's circle: the evidence for the premises warranting study-target comparability must not support another argument in which the conclusion being extrapolated is warranted independently of the study results. In this section, I examine and reject EBM's solution to the problem of extrapolation, simple extrapolation (in subsection 7.4.1). I then propose a new solution to the evidential problem of extrapolation that harnesses structural analogies (in subsection 7.4.2).

### 7.4.1 Simple extrapolation

In 2011 in the journal *Lancet*, Cartwright wrote about the problem of extrapolation in a study, or in her words, providing justification that the intervention "will work for us" in our target context based partly on the finding that it "works somewhere." In a letter response, Mark Petticrew and Iain Chalmers argued, "Cartwright, like many others, conceptualises the challenge as being to demonstrate that the characteristics and circumstances of the research are sufficiently similar to those to which extrapolation is being contemplated. But why should the challenge be conceptualised that way round? Why not instead ask 'Are there any good reasons to believe that the research is not relevant to us, that "it won't work for us"?' If there are not, and considering the undesirable alternative ways of reaching a decision, the default position should be that the result should be regarded as applicable" (2011, 1696). Petticrew and Chalmers are not alone among EBM experts in advocating such an approach.[23] Piet Post and colleagues (2013) did a systematic review of the medical literature to identify approaches to generalizing trial results. Nine of fifteen results that met their inclusion criteria recommend the strategy of Petticrew and Chalmers, the method of "evaluating whether there are compelling reasons why the relative effect found in trial results should not be applied to the patient group" (Post et al. 2013, 639). Among them was

[23] Other advocates include Dans et al. (1998); Sheldon et al. (1998); N. Weiss (2006); Straus et al. (2011); Guyatt, Oxman, Kunz et al. (2011); Guyatt et al. (2015).

the popular GRADE system for rating evidence (Guyatt et al. 2008). GRADE endorses the approach for assessing study population–target population comparability, which is one dimension of the "directness" of the effect estimate in GRADE (Guyatt, Oxman, Kunz, et al. 2011). Summarizing, Post et al. argue that we should be "accepting that results of randomized trials apply to wide populations unless there is a compelling reason to believe the results would differ substantially as a function of particular characteristics of those patients" (2013, 641–642).

We can call this approach "simple extrapolation":[24] extrapolate the relative effect size from the trial unless there is a compelling reason not to extrapolate. By "the relative effect size," simple extrapolation proponents mean the RRR or the relative risk (RR; RR = 1 – RRR). Compelling reasons are essentially reasons to believe that the effect size will not be the same in the particular target population due to characteristics of target patients or the target setting. Compelling reasons are uncommonly strong reasons that will not be found in most cases. Applied to the efavirenz trials, extrapolators had a compelling reason to think that the effect size would differ in Zimbabwe but seemingly ignored it (Park et al. 2023), perhaps failing to properly implement simple extrapolation, which would have led them to the correct conclusion that the effect size would not generalize. Applied to the initial trials of endarterectomy for stroke prevention, extrapolators might *not* have had a compelling reason to think that the effect size would differ in certain other populations at the time, lacking complete data on its sensitivity to the degree of carotid narrowing, sex differences in atherosclerotic plaque pathology, and the vulnerability of the arteries to stroke in the initial two weeks following a previous stroke event. They might have correctly applied simple extrapolation and accepted the wrong conclusion.

Simple extrapolation is a solution to the evidential problem of extrapolation because it tells an extrapolator how to use evidence when extrapolating: presume that the RRR is generalizable unless you have strong evidence that says otherwise. As a solution to the evidential problem, it can be paired with an extrapolation argument in which the study result being extrapolated is the effect size, such as the Generalization argument from section 7.3.3. Simple extrapolation avoids the evidential extrapolator's circle

[24] Steel (2008) analyzes the same method used in toxicology to extrapolate harms from animals to humans, calling it "simple induction." Jacob Stegenga (2015) criticizes the EBM version, referring to it as "simple extrapolation, unless" because it recommends extrapolating *unless* there is a compelling reason.

because it does not require confirming evidence that establishes the sufficient similarity assumption, and the evidential extrapolator's circle is violated only when an extrapolator has such evidence. Simple extrapolation gets around the extrapolator's circle by not requiring the extrapolator to have evidence in favor of the sufficient similarity assumption at all: they are licensed to accept the assumption in the absence of disconfirming evidence. But how can simple extrapolation get away with this?

Though many simple extrapolation adherents are silent on why they think it is a reasonable approach, some explicitly justify it with the assumption that in most cases, as a matter of fact, the relative treatment effect will be the same in the target and the study.[25] For example, Post et al. claim that "in general, relative effect of treatment on benefit outcomes seldom differs to an important extent across subgroups of patients" (2013, 638). I call this assumption "general generalizability" (GG). It appears to function as a probabilistic warrant for extrapolation: the assumption that most RRRs are generalizable to most target populations is meant to confer high probability on the conclusion that *this* RRR is generalizable to *this* target population. Thus, to defeat the conclusion, one would need a compelling reason to believe otherwise. As David Sackett and colleagues observe, "[T]his (constancy of RR) is a big assumption" (1997, 170). Big assumptions demand strong justification. In most of the rest of this subsection, I will argue that strong justification is lacking. The evidence offered in support of GG is wanting, and, in fact, there are good reasons to be skeptical of GG.

It is commonly claimed in the EBM and methodological literatures that relative measures like the RRR are more likely to be generalizable compared to absolute measures like the ARR, including across different levels of baseline risk (though not everyone agrees).[26] One could question this claim. There are scenarios—described in section 7.3.3—in which the opposite will be true. Regardless, GG is not the assumption that relative measures are *more often* generalizable than absolute measures; it is the assumption that they are *generally* generalizable. It could be that relative measures are more often generalizable but not often generalizable.

[25] For instance, see Dans et al. (1998); N. Weiss (2006); Guyatt, Oxman, Kunz et al. (2011); Post et al. (2013); Guyatt et al. (2015); R. Collins et al. (2020).

[26] See Colnet et al. (2024) on sources claiming that one measure or another is more likely to be generalizable, along with a formal analysis. See Poole et al. (2015) on sources claiming the ARR or risk difference is a more heterogeneous measure, along with a critique of the evidence cited.

Simple extrapolation proponents don't always offer support for GG even when it is explicitly invoked. But when they do, that support sometimes takes the form of empirical research similar to that cited by Post and colleagues (2013).[27] The authors reference meta-research showing low variation in the relative risk among RCTs included in a meta-analysis, including low variation across levels of untreated risk (Furukawa et al. 2002; Deeks 2002). In one study, the relative risk measured in an RCT was not significantly different (statistically speaking) from the pooled meta-analytic relative risk around 90% of the time (Furukawa et al. 2002). Concerns have been raised about the statistical power to detect differences in these sorts of analyses (Poole et al. 2015). Furthermore, the RCTs included in a meta-analysis may be unusually similar for several reasons. First, RCTs included in a meta-analysis are often similar in their eligibility criteria and trial protocol. Moreover, these eligibility criteria are often restrictive in ways that might matter for the effect size, as discussed below. In contrast, patients and the care they receive can vary widely in routine practice. Finally, a meta-analysis might not be run in the first place if substantial treatment effect heterogeneity is suspected or observed in RCTs. All of these considerations suggest that the RCTs included in a meta-analysis might constitute an unusually homogeneous sample of target populations in which to test GG.

A better empirical test of GG would require a representative sample of interventions and target populations/target settings and would examine whether the relative treatment effect for each intervention differed substantially among diverse subgroups representative of healthcare. Lacking such an ambitious empirical undertaking, we must instead search for theoretical justification for GG.

In section 7.3.3, mechanism-cofactor analysis showed that differences in a population's distribution of intervention cofactors and disease mechanisms can result in differences in the RRR. The relative treatment effect will be similar among populations only if there is little variation in mechanisms and cofactors among individuals, or if there is variation but in aggregate the population-level distribution of mechanisms and cofactors is similar in just the right ways. On the contrary, I will argue that there are good general reasons to expect variation across individuals, and that this variation will often aggregate to produce differences in the relative treatment effect.

[27] Spiegelman and VanderWeele (2017) provide anecdotal evidence for the relative constancy of relative measures in public health.

First, there are good general reasons to expect variation in disease mechanisms and intervention cofactors across individuals. According to epidemiology's prevailing multifactorial theory, disease outcomes are caused by the interaction of many risk factors, including genetic, behavioral, environmental, and social causes. These causal risk factors commonly vary among individuals, producing variation in disease mechanisms. The best evidence for variation in the mechanisms responsible for common disease outcomes such as cardiovascular events and cancers is provided by the observation that there are vast differences in the occurrence and timing of these outcomes among individuals, as well as vast differences in the absolute risk of these outcomes across populations. Moreover, genetic, behavioral, environmental, and social causes are also among an intervention's cofactors, including factors involved in the absorption, distribution, metabolism, and elimination of drugs (Lai et al. 2022), patient adherence to the intervention, concurrent diseases and interventions, and the availability and quality of healthcare. Richard Kravitz and colleagues (2004) provide several examples of individual treatment effect heterogeneity produced by genetic, behavioral, and contextual factors.

Variation in mechanisms and cofactors across individuals will often aggregate to produce differences in the relative treatment effect across populations. Examples examined in previous sections included efavirenz's neuropsychiatric side effects across countries, trastuzumab's effectiveness in preventing death across older versus younger patients, and the variable effect of endarterectomy across groups that differ in their degree of carotid stenosis, sex, and time since their most recent stroke event. Further examples of relative treatment effect modification are compiled by others (Kent et al. 2018; Rothwell et al. 2005). Figure 7.3 illustrates how differences in the individual-level effects of an intervention can sum to create differences in the relative risk across target populations. In fact, there are more ways that the individuals in Figure 7.3 can be grouped into populations to produce divergent relative risks than there are ways they can be grouped to create similar relative risks.

In the setting of individual variation in mechanisms and cofactors, the relative treatment effect will be similar in only certain scenarios, examined in section 7.3.3. According to the sufficient similarity assumption, the effect size will be similar in the target population if the intervention's cofactors and the disease's mechanisms are distributed similarly in the study population (and similar causal principles are true of both settings). Otherwise, differences

in their distribution must "cancel." (For instance, an increase in disruption cofactors must be balanced by an increase in harm cofactors.) Or it must be that only targeted mechanisms are distributed differently, and other special circumstances must also obtain. Unfortunately, there are good reasons to worry that, in general, these scenarios may not be common.

In particular, there is good reason to worry that the RRR measured in a clinical trial may not be widely generalizable because clinical trials are often atypical; they tend to differ from routine medicine in ways that can influence the RRR. For instance, given the changing characteristics of patients in community and hospital practice, target populations increasingly contain many older patients, patients with multiple diseases, and patients taking multiple medications, yet RCTs often exclude these very patients (Van Spall et al. 2007; Kennedy-Martin et al. 2015; Tan et al. 2022). Age-related physiological changes, disease-related pathophysiological changes, and medications commonly are among an intervention's cofactors and can also spur disease mechanisms. For instance, age-related cardiovascular changes, degree of carotid stenosis and carotid plaque morphology, concurrent atrial fibrillation, and concurrent statin use can all influence the effectiveness of endarterectomy (Rothwell et al. 2003, 2004, 2005). Moreover, the untreated absolute risk will often differ in the target compared to the trial (Rothwell et al. 2005), and the untreated absolute risk is determined by the distribution of targeted and non-targeted disease mechanisms. Finally, patient care often differs between study settings and typical care settings, sometimes in ways that are shown to influence treatment effectiveness (Rothwell 2005a). These dissimilarities between typical trials and routine medicine provide good reason to doubt GG.

In summary, GG is a problematic assumption, and simple extrapolation is a problematic solution to the evidential problem of extrapolation for depending on it.[28] We should worry that the relative effect size, measured in a clinical trial, will fail to generalize to a target population due to differing distributions of cofactors and mechanisms. The idea of the "golden risk ratio," that each intervention's RRR is a universal constant, is a dangerous myth. In general, rather than assuming GG and searching for compelling evidence against the generalizability of the relative effect size, extrapolators

[28] See Fuller (2013a, 2021b) for further analysis of justifications for GG. If we strip simple extrapolation of GG, it is a bare argument from ignorance: there is absence of evidence (of non-generalizability of the relative treatment effect); therefore, there is evidence of absence (of non-generalizability). I argue that neither epistemic nor practical considerations can rescue simple extrapolation from this fallacy (Fuller 2021b).

should seek evidence for and against generalizability and weigh all this evidence in deciding whether to extrapolate the effect size.

Despite simple extrapolation's shortcomings as a general approach, its use might sometimes be warranted in a restricted domain, for instance, in the use of an intervention for which we know from empirical evidence or biological theory that relative effectiveness varies little among individuals or populations, or when the potential benefits of extrapolating and intervening almost always outweigh the potential harms. However, even when the relative treatment effect is generalizable, it is less informative for treatment decision-making than the absolute treatment effect (see Chapter 6 as well as Sprenger and Stegenga 2017). Thus a decision-maker should either extrapolate the absolute treatment effect from therapeutic research instead, or they should use the approach examined in section 7.3.3 and extrapolate the relative treatment effect from therapeutic research, the untreated absolute risk from prognostic research, and use them to calculate the absolute treatment effect in the target. In other words, even when the relative effect is generally generalizable within a restricted domain, the extrapolator should still refrain from using simple extrapolation when they can extrapolate the absolute effect instead or when they are unable to calculate the absolute effect from the relative effect because they lack good prognostic research or because the untreated absolute risk estimated in that research is not generalizable to the target.

Given its shortcomings, we should look beyond simple extrapolation for a solution to the evidential problem of extrapolation.

### 7.4.2 Structural Analogies

The challenge of the evidential version of the extrapolator's circle is the puzzle of how an extrapolator could establish using evidence that study and target populations are causally comparable without that evidence establishing a prediction about the target population directly. How can one know that two things are identical without knowing their identity? Though the problem seems paradoxical, it is not always insurmountable. For instance, we can know that two people are identical twins without knowing much about them other than that they developed from the same fertilized egg. When it comes to causal identity, my solution is to draw structural analogies between the study and target that evidence functional

analogies. This solution to the evidential problem overcomes the evidential extrapolator's circle.

All three of the solutions to the logical problem explored in section 7.3 involved comparing the study and target for the presence or absence of mechanisms and causes. Broadly speaking, what determines whether a factor is a cause of an outcome, whether it plays a certain causal or functional role, is its structure. Because structure determines causal role, structural analogies can provide a basis for causal analogies. For medical interventions such as medications and surgeries, the relevant kinds of structure are given to us by biomedical science. Rather than providing a general account of "structure" to ground my solution, I will explore different kinds of structure studied in the biomedical sciences through several examples.

The kinds of structure-function relationships studied in the biomedical sciences are numerous and varied. For instance, muscles sharing the same mass and arrangement of muscle fibers generate the same force given the same stimulus (etc.); enzymes sharing the same three-dimensional shape catalyze the same reaction given the same substrates and coenzymes (etc.). However, these effects might not be predictable from knowledge of the structure alone. One might not be able to predict the force by studying the muscle's mass and arrangement, nor predict the enzymatic products by studying the enzyme's shape. This lack of predictability can successfully block the extrapolator's circle. To use the language of mechanism-cofactor analysis, if a biological factor in the target is structurally similar to a cofactor for the intervention in the study, an extrapolator might be able to infer that the biological factor plays the same causal role as the cofactor in the study, and thus that it is a cofactor in the target. However, without making the structural analogy, the extrapolator might not be able to predict the effects of the biological factor in the target.

This general strategy could work for both macroscopic and microscopic structural analogies. As an illustration, macroscopic atherosclerotic plaque diameter and microscopic plaque morphology are both components of sufficient causes of stroke, as evidenced by the fact that they influence both stroke risk and the effectiveness of endarterectomy (Rothwell et al. 2003, 2004). If the extrapolator had evidence that the diameter of the plaque was similar in the target as in the study, or that plaque morphology was similar, they could infer that, ceteris paribus, these factors were performing the same causal role in the target without understanding fully the biology of plaque rupture. If they understood more fully how plaque diameter and morphology influence

plaque rupture, they might be able to infer that a certain plaque diameter or morphology would play a certain causal role in the mechanism without making a structural analogy, and the extrapolator's circle would be sprung. Rather, the extrapolator might know that plaque diameter plays a role in plaque rupture, but not the threshold diameter above which it becomes relevant; nonetheless, they can infer that similar plaque diameters have similar relevance (given other relevant causal similarities).

Structural biology is a discipline that studies the structure and function of biological macromolecules, including proteins. Two proteins that have the same structure (referring to properties like shape and distribution of electrostatic charge) often perform the same causal role by binding to the same ligands with the same binding affinity (Lean 2020) in the same circumstances. For example, the reason that the defective CYP2B6 variant fails to play the same causal role as the wild-type CYP2B6 enzyme is due to a point mutation in its gene: a guanine-to-thymine nucleotide base pair substitution. This results in a glutamine-to-histidine amino acid substitution in the CYP2B6 protein (Nyakutira et al. 2008). This change to the amino acid sequence or "primary structure" of the protein alters its shape enough so that it no longer metabolizes efavirenz efficiently. Thus the change in structure is enough to change the protein's function, so that it no longer plays the same causal role of metabolizing efavirenz (at least not with sufficient efficiency to avoid a high blood concentration of the drug). Genetic testing can determine which structural variant a patient has, and thus what cause the patient has; two patients with the same genetic mutation can be presumed to have a structurally similar enzyme that plays a similar causal role (other things being equal). This structural analogy does not imply the ability to predict the causal role of the protein from the structure alone.

Structural analogies can be used not only to establish that two factors function as the same intervention cofactor but also to establish that two factors function as the same intervention, meaning an intervention that interacts the same way with the same cofactors. Take the example of trastuzumab (tradename Herceptin), which is a targeted treatment for HER2-overexpressing breast cancer (Hudis 2007; Vu and Claret 2012; Hayes 2019). An aggressive subtype of breast cancer involves amplification and overexpression of the gene responsible for human epidermal growth factor receptor 2 (HER2) in breast cancer cells. HER2 is a transmembrane protein with tyrosine kinase activity. It is activated by dimerizing with other HER2

receptors or with related epidermal growth factor receptors. Meanwhile, trastuzumab is a recombinant monoclonal antibody. Like all antibodies, it has both Fc and Fab regions. The Fc region binds to cell surface receptors, while the Fab region binds to antigens and is antigen specific.

Trastuzumab works by binding to HER2's extracellular domain at its Fab region. It is theorized that this binding inhibits uncontrolled cell growth and survival in several ways (Vu and Claret 2012). One way is by blocking dimerization of HER2. Another way is by recruiting immune cells, especially natural killer cells, that bind to trastuzumab's Fc region and destroy the cancer cell. Trastuzumab has been shown in clinical trials to improve survival when used after breast cancer surgery (adjuvant treatment) as well as in patients with metastatic breast cancer (Wolff et al. 2018). The structure of Herceptin explains its causal role in breast cancer treatment. The binding of Herceptin to both HER2 and to immune effector cells is highly specific, dependent on the structure of its Fab and Fc regions, respectively. Another antibody with structurally similar binding regions (such as a generic brand of trastuzumab made by a different manufacturer) can be expected to bind to the same protein ligands given similar background conditions. An extrapolator might not be able to predict this binding affinity just by studying the protein structure, but they might still be able to predict that the two brands of trastuzumab would have similar binding affinity for HER2 and for immune effector cells because of their similar structure.

Two proteins with the same structure will have the same complete binding specificity[29] profile, meaning that they will have equal affinity for binding protein A, equal affinity for binding protein B, and so on (ceteris paribus), for all proteins. In other words, they will have all the same effects and side effects. In extrapolating the effects of one HER2 inhibitor to a population that uses a different kind of HER2 inhibitor (such as pertuzumab; Hayes 2019), an extrapolator needs to predict not only that the second kind of HER2 inhibitor will bind to HER2 with the same affinity but that it will bind to off-target sites with the same affinity, that it will bind to proteins that metabolize or eliminate the inhibitor with the same affinity, and so on, so that it will have the same toxicities as the first inhibitor and not only the same intended effects.

My structural analogies proposal is a solution to the evidential problem of extrapolation because it suggests what kind of evidence would support those

[29] Oliver Lean (2020) develops a causal interpretation of binding specificity through an analysis of drug design.

premises in an extrapolation argument that warrant study-target comparability. Its answer is that this evidence can consist in a structural comparison between (i) the intervention and other component causes of the outcome in the study population known to play certain causal roles and (ii) factors in the target population. How we make a structural comparison will vary. For instance, a similar atherosclerotic plaque structure can be established through medical imaging that reveals plaque diameter and morphology. Meanwhile, a similar protein structure will likely be established less directly through genetic testing or protein assays that establish the identity of a protein that has a structure known from research in structural biology. Evidence of structural similarity can be used to establish causal similarity. This solution to the evidential problem avoids the evidential extrapolator's circle so long as knowledge of the structure alone is not enough to reveal the factor's causal role.

My solution to the evidential problem can be paired with any of the three solutions to the logical problem examined earlier. For instance, an extrapolator can use it to extrapolate causal relevance using my modified version of comparative process tracing (section 7.3.1). To warrant study-target comparability in this extrapolation argument, the extrapolator must establish that there are factors in the target patient structurally similar to causes that interacted with the intervention in the study patients. These structural analogies need only be strong enough to establish that those factors play the same qualitative causal role of interacting with the intervention to raise (or lower) the probability of the outcome by some amount. In the efavirenz example, the structure of the CYP2B6 enzyme must be similar enough to the inefficient variant to raise the probability of efavirenz's side effects.

If we wish to instead predict the effect size using transport formulas (section 7.3.2) or mechanism-cofactor analysis (section 7.3.3), we require additional evidence of the distribution of the relevant disease mechanisms and intervention cofactors in both the study and the target. However, recall that the relevant extrapolation arguments must now assume that similar *quantitative* causal principles governing the intervention and its outcome are true of both settings. The structural similarities must therefore be stronger, enough to establish not just that factors play the same qualitative causal role but that they interact to produce the same probability of the outcome. For instance, a medication's cofactors will usually include liver enzymes like CYP2B6 that metabolize the medication, either activating it, inactivating it, or facilitating its elimination from the body. In the case of the CYP2B6 variant we have been discussing, its ability to metabolize efavirenz was so

reduced, especially in people homozygous for the variant allele, that we considered it to be a cause distinct from the wild-type variant. But we could instead consider these variants to be different subtypes of the same enzyme (Nyakutira et al. 2008). People who possess the lower efficiency subtype of an enzyme are often called "slow metabolizers" (Zangger and Schwab 2013). The effect of enzymatic subtypes is qualitatively similar—they produce the same products from the same substrates—but quantitatively different owing to their different kinetics. In predicting the effect size, structural analogies must establish quantitative causal similarity.

Due to the complexity of protein structure-function relationships, it will often be difficult to determine—using our biological understanding of how structure is related to function—when two factors are structurally similar enough that they will have the same qualitative effect but dissimilar enough that they will have a different quantitative effect; making this determination may also require the level of detailed biological understanding that invites the extrapolator's circle. For instance, one might think that two proteins that share a similar but nonidentical amino acid sequence have qualitatively similar effects, while two proteins that share the exact same amino acid sequence have quantitatively similar effects. However, this general principle will often betray us because even large differences in amino acid sequence can be completely inert, while even small differences can have qualitative consequences (Teng et al. 2010), as in the example of CYP2B6's functionally influential single amino acid substitution. Even proteins with identical amino acid sequences may be functionally distinct because the amino acid sequence (primary structure) is not the only determinant of protein structure. Covalent modifications, non-covalent binding of other molecules, and other interactions also influence functionally relevant protein structure. Small changes in structure can result in substantial differences in effect that are difficult to predict in advance from structure alone due to the complexity of protein structure-function relationships (Mitchell and Gronenborn 2017). Thus an extrapolator might be able to make a secure structural analogy only when the relevant structure is identical or nearly identical, in which case their analogy might be strong enough to secure both qualitative and quantitative causal similarity.

In summary, while simple extrapolation attempts to solve the evidential problem of extrapolation by forgoing positive evidence of study-target comparability and extrapolating in the absence of evidence against comparability, its presumption of the general generalizability of the relative effect

size is unjustified. On the other hand, my structural analogies solution to the evidential problem uses positive evidence of structural similarity between factors in the target and causes in the study. It evades the extrapolator's circle because structure-function relationships in biology are often difficult to predict directly. But because similar structure often establishes similar causal role, we can extrapolate the effect of structurally similar biological factors once we have measured their effect in an empirical study such as a clinical trial. My structural analogies solution can be paired with any of solutions to the logical problem investigated earlier to provide an adequate (if partial) solution to the problems of extrapolation.

## 7.5. External Validity of Clinical Trials and the Practical Problem of Extrapolation

What are the upshots of my solutions to the logical and evidential problems of extrapolation for the ominous problem of the external validity of clinical trials? One part of the problem is the concern that many of medicine's clinical trials do not have external validity for typical target populations, which amounts to the concern that our extrapolation argument will be unsound. This concern is well motivated when extrapolating the effect size because, as we saw, the relevant sufficient similarity assumption is demanding, requiring a similar distribution of disease mechanisms and intervention cofactors that can each consist of myriad component causes. Moreover, we also saw that we lack good universal reasons to assume that the RRR, a common measure of the effect size, is generally generalizable from trials to target populations. However, a different solution to the logical problem that utilizes transport formulas sometimes allows an extrapolator to predict the effect size in the target even when that effect size differs from the effect size in the overall clinical trial population.[30]

Because causal relevance requires a less demanding sufficient similarity assumption, the external validity of this conclusion will often be more widespread compared to the effect size. That is to say that there will more often be some patients in the target and in the trial who respond in a qualitatively similar way to the intervention. The trade-off is that causal relevance is a

[30] Steel (2008), Stuart et al. (2010), Westreich et al. (2017), Deaton and Cartwright (2018), and Khoswrowi (2019) also make suggestions about reweighting the effect size when extrapolating.

correspondingly more modest conclusion, too modest to ground many clinical decisions. It is often not enough to know that an intervention makes some causal difference for some patients in a target population, for such information often underdetermines whether the intervention should be used. Harm-benefit analysis in therapy involves weighing the harms and benefits of an intervention, and the weights of an intervention's benefits and harms depend partly on the magnitude of benefit and harm.

The other component of the problem of the external validity of clinical trials is the general problem of how to assess external validity using arguments and evidence, which I have solved in principle with my solutions to the two problems of extrapolation. But how feasible are these solutions in practice? My solutions to both problems of extrapolation instruct the extrapolator to represent the relevant causes and mechanisms (using, for instance, causal graphs and causal pies). This allows them to identify the factors (e.g., metabolic enzymes) that causally interact with the intervention, and sometimes also at which mechanism stages. They can then determine using structural analogies if similar causal factors function in the study and the target. This entire exercise requires a level of causal-mechanistic understanding that an extrapolator may not possess. In fact, a virtue of running a randomized trial is that it requires little knowledge of the relevant causes and mechanisms to make a causal inference and estimate the effect size.[31] For instance, it is commonly argued that a trialist need not worry about what causes could confound the trial's analysis because randomization prevents those causes from doing so (though see Deaton and Cartwright 2018 on limitations to this argument).

Thus one might worry that my solutions to the problems of extrapolation imply that we can rarely (if ever) extrapolate from clinical trials, a conclusion that could be seen as either unfortunate or implausible. The worry that this implication is unfortunate is of course not a good objection unless we assume that scientific inference is always easy and that we always get what we want. The worry that this implication is implausible is more serious, but I reject the implication because it is based on a misunderstanding of the problems of extrapolation that I have set out to solve. The logical and evidential problems of extrapolation are about what arguments and what evidence are relevant in extrapolating. They are not the problem of when we should extrapolate,

[31] To use Cartwright's (2007) terminology, a randomized trial is "self-validating." However, Brendan Clarke and colleagues (2014) argue that designing a clinical trial *does* require some evidence of mechanisms.

meaning when we should accept that our argument is sound or that our evidence is strong enough to act.

There is a distinct problem of extrapolation, a practical problem of extrapolation, which is the problem of when to extrapolate from a clinical trial or other study. This problem is not so much epistemological but, as its name suggests, practical. A solution must be sensitive to the practical consequences of extrapolating: what an extrapolator stands to gain and lose if they are right and if they are wrong in accepting or rejecting the conclusion (Douglas 2000; Djulbegovic and Hozo 2007). It also depends on the ethics of extrapolating, the ethical values or principles that influence when physicians or policymakers should accept that a study result is generalizable. The problem of when to extrapolate depends on these consequences and ethics because accepting that a study result is generalizable in medicine typically has implications for practice or policy. Extrapolating a trial result could lead a physician to recommend an intervention (or not), which will have practical consequences for the patient. It could also lead a regulator to approve an intervention (or not), an insurer to fund an intervention (or not), or a policymaker to mandate an intervention (or not).

In deciding whether to extrapolate from a study, an extrapolator typically needs in hand solutions to all three problems, the logical, evidential, and practical. A solution to the practical problem does not render solutions to the logical and evidential problems needless because an extrapolator needs to know not only the consequences of extrapolating and intervening but also the probability or plausibility of those consequences obtaining, which requires arguments and evidence. I have provided solutions only to the logical and evidential problems in this chapter because those solutions are highly specific to the problem of extrapolation, asking us to formulate a particular kind of argument and search for particular kinds of evidence. In contrast, a solution to the practical problem of extrapolation is likely to be more general. In other words, it would likely rely on theories of decision-making or ethical theories that apply just as well beyond the context of extrapolating from a clinical trial.

My solutions to the logical and evidential problems, and the causal-mechanistic understanding they require, have the implausible implication that an extrapolator can rarely extrapolate from clinical trials on only two conditions: (1) we are typically not highly certain that our causal-mechanistic understanding relevant to extrapolating is sound and detailed, and (2) our solution to the practical problem of extrapolation insists that

we extrapolate only when we are highly certain that our causal-mechanistic understanding is sound and detailed. Because there are situations in which we have much more to gain from extrapolating and using an intervention than we have to lose, such a solution to the practical problem would be unacceptable.[32] Thus my solutions to the logical and evidential problems (fortunately) do not carry the implication that we can rarely extrapolate from clinical trials.[33]

Jeremy Howick and colleagues (2013, 288) argue that "mechanisms approaches" to the problem of extrapolation are "rarely successful." They provide four objections to mechanisms approaches: "First, our knowledge of underlying mechanisms is often mistaken or incomplete. Secondly, knowledge of mechanisms often cannot be justifiably extrapolated outside the tightly controlled laboratory situations in which such knowledge is usually produced. Thirdly, mechanisms can behave paradoxically. Finally, using mechanistic knowledge does not overcome what Dan Steel calls 'the extrapolator's circle'" (276). I have already resolved their fourth objection in sections 7.3 and 7.4. The first three objections concern deficiencies in our understanding of the relevant mechanisms. (The third objection is really only a problem if we do not know in advance when a mechanism will behave paradoxically, having an opposite effect.) My answer to their first three objections here is that these problems should reduce our certainty in our causal-mechanistic understanding, but on any plausible solution to the practical problem of extrapolation we can still use mechanisms approaches to extrapolate when we are uncertain.

[32] Analogous reasoning can be wielded against simple extrapolation's presumption that we should generally extrapolate from clinical trials. Namely, simple extrapolation is insensitive to the consequences and ethics of extrapolating (Fuller 2021b), and thus implies an inadequate solution to the practical problem of extrapolation. For simple extrapolation, the problem is the reverse of the one I've been describing: simple extrapolation may lead us to extrapolate from studies *more often* than we should. This could in turn lead to overtreatment, which is a much-maligned problem in contemporary medicine (Moynihan and Smith 2002).

[33] In addition to knowing what assumptions need support, what evidence can support the assumptions, and when to extrapolate under conditions of uncertainty, we also need to assess how uncertain we are given our assumptions and evidence. Donal Khoswrowi (2023) proposes a "support graph approach" for evaluating uncertainty when extrapolating.

## 7.6. Conclusion

The problems of extrapolating from clinical trials are twofold: a logical problem and an evidential problem. The logical problem demands a good extrapolation argument, while the evidential problem demands to know what evidence could support those premises in the argument warranting study-target comparability. I developed my own solution to the logical problem, mechanism-cofactor analysis, from my analysis of preventive and curative medical interventions, as well as my own solution to the evidential problem that utilizes structural analogies. An adequate solution to the problems of extrapolation must address both problems, and even a partial solution must contend with the extrapolator's circle. The extrapolator's circle is violated when the premise referring to the study's results is dispensable or when the evidence for the premise referring to study-target comparability can establish the conclusion independently of the extrapolation argument and the study results. Solutions to the logical problem vary depending on the study result being extrapolated but generally rely on an analysis of relevant mechanisms and causes. However, existing proposals do not generally tackle the evidential problem. An exception is simple extrapolation, which avoids the extrapolator's circle but by making the unjustified assumption that the relative treatment effect is generally generalizable. Thus I proposed a new solution to the evidential problem relying on structural analogies: the extrapolator can use structural similarities between factors in the study and the target to establish causal similarities. An upshot of my analysis is that assessing the external validity of clinical trials depends on having causal-mechanistic understanding, but this does not imply that we can rarely extrapolate from trials because any plausible solution to the practical problem of extrapolation (when to extrapolate, given the practical and ethical context) should allow us to extrapolate in conditions of uncertainty.

In the next chapter, I return to the source—medical evidence—and ask how we should calibrate our skepticism about medical evidence to empirically demonstrated research biases.

# 8

# Therapeutic Skepticism


## Summary

What epistemic attitude should we adopt toward the results of contemporary therapeutic research and toward contemporary therapies, especially given ubiquitous meta-research evidence of bias in the medical literature? I defend Medical Meta-Evidentialism, which recommends that we generally lower our confidence in therapeutic study results and become less optimistic about therapeutic effectiveness having learned of this meta-research evidence. Meta-Evidentialism provides a better response to the problem of biased and misleading therapeutic evidence compared to Medical Nihilism, which instead asks us to adopt a low confidence in contemporary medical interventions.

Throughout the history of modern medicine, therapeutic skepticism has waxed and waned, driven by forces inside and outside of medicine. Today much medical skepticism is driven by empirical findings suggesting that therapeutic research is biased or misleading, including findings from meta-research, or research on medical research. Bias occurs when evidence misleads us about the soundness of our medical inferences. While many commentators have proposed ways to mitigate the influence of methodological and social biases in future research, we also face the challenge of how to account for the influence of these problems on our current evidence. One response is Medical Nihilism. I advance three objections to Medical Nihilism: the broad-brush objection, the multiple measures of effectiveness objection, and the vague-probabilities objection. These objections motivate my alternative, Medical Meta-Evidentialism. Meta-Evidentialism avoids the objections to Medical Nihilism and is bolstered by an argument from meta-research evidence, which justifies taking meta-research evidence to be medical evidence as well as using this medical evidence in therapeutic evaluation. Medical Meta-Evidentialism provides a viable path forward between the thorns of unbridled therapeutic optimism and the quicksand of excessive nihilistic despair.



*The New Modern Medicine*. Jonathan Fuller, Oxford University Press. © Oxford University Press 2025.
DOI: 10.1093/9780190066178.003.0009

## 8.1. Cycles of Therapeutic Skepticism

In Part 3 of this book, I have explored evidence and inference in the new modern medicine. In Chapter 5, I examined evidence-based medicine's (EBM's) measurement view of medical evidence and its standard model of medical prediction, which involve measuring and extrapolating population-level risks and effect sizes from epidemiological studies that then inform patient-level probabilities. In Chapter 6, I interpreted population-level effects of preventive and curative medical interventions as well as patient-level probabilities inferred from epidemiological evidence. In Chapter 7, I interrogated the problems of extrapolating treatment effects from clinical trials. Throughout these chapters, I have often leaned on an argument theory of medical evidence, in which medical evidence supports premises in arguments for medical hypotheses. In this final chapter of Part 3, the argument theory helps us understand how meta-research findings can provide meta-evidence for evaluating the effectiveness of medical interventions, as well as what it means when medical evidence is misleading or biased.

There have always been therapeutic skeptics—doctors and patients who have taken a skeptical stance toward medicine, ranging from the questioning to the downright nihilistic. Over time, mainstream medicine's self-opinion has swung between optimism and pessimism regarding its therapies (Greene et al. 2012). For instance, in the first half of the nineteenth century, Paris-trained physicians practicing their own brand of scientific medicine imported a therapeutic nihilism into America. Armed with Louis's "numerical method," Paris hospital physicians had begun comparing large numbers of people on different therapeutic regimens and often found that aggressive treatments did not work as well as previously thought. This crisis of confidence was resolved only after the rise of the germ theory in the second half of the century and the delivery of effective treatments like insulin for diabetes and chemotherapies for infections in the first half of the next century. In the middle of the twentieth century, a pharmaceutical revolution delivered several new classes of drugs, including antibiotics, psychiatric medications, and preventive therapeutics for noncommunicable diseases (Greene et al. 2012; Le Fanu 2012). These discoveries were triumphs of "laboratory medicine," in which laboratory sciences had paramount influence on medical practice (Chapter 1).

Then in the 1970s, a resurgent medical nihilism was fueled by social and historical critiques of modern medicine as well as broader critiques of

Western institutions. In *Medical Nemesis*, Ivan Illich declared that the medical establishment had "become a major threat to health" (1976, 11), arguing that iatrogenesis, in which medicine does harm, was widespread in its traditional guise of noxious treatments (as well as in the form of medicalization of health and diminishment of the individual's autonomy over health). Meanwhile, Thomas McKeown (1976a, 1976b) argued that for common infections like tuberculosis, declines in mortality over the previous two hundred years—a feature of the classic "epidemiologic transition" discussed in Chapter 1—had little to do with medical interventions like vaccines or antibiotics and much more to do with improving socioeconomic conditions, especially improvements in nutrition. Looking back, some have argued that this was a period during which the midcentury swell of therapeutic innovation dried up, resulting in a dearth of truly effective new treatments (Le Fanu 2012), as medicine left behind its golden age (Varga 2022).[1]

While many of Illich's and McKeown's ideas were criticized by scholars and largely dismissed or ignored by mainstream medicine (Colgrove 2002), those of another skeptic of the era helped spur a more radical medical response. Epidemiologist Archibold Cochrane's 1972 monograph *Effectiveness and Efficiency* lamented the unclear and dubious effectiveness of many medical interventions and the resources being wasted on them. Cochrane's two nemeses were tuberculosis and research biases (errors to which poor quality epidemiological studies were prone). He had studied under Bradford Hill, and in the 1948 randomized controlled trial (RCT) of streptomycin for tuberculosis, which implemented Hill's vision for randomized clinical experiments, Cochrane saw an antidote for both scourges (Daly 2005). He argued for medical RCTs to root out bias and for reviews of the medical literature to catalog what worked and what didn't. These ideas would later inspire the Cochrane Collaboration, a network of collaborators running systematic reviews of the medical literature (including systematic reviews of RCTs) that was founded in 1993 in Cochrane's name.

Cochrane's ideas also became central to the EBM movement around the same time (EBM Working Group 1992). EBM focused on using clinical epidemiology research, especially RCTs and systematic reviews, to generate therapeutic evidence that is as free of bias as possible. EBM also implemented

[1] Meanwhile, medicine and manufacturers expanded the use of treatments that were effective for some patients to wider swaths of people, some of whom might not benefit and in fact might be harmed. This concern is a central one within the contemporary "too much medicine" movement, a doctor-driven crusade to curb overtreatment (Moynihan and Smith 2002; Le Fanu 2019).

standards for evaluating evidence; for example, the "Cochrane risk-of-bias" tool became a standard instrument for systematic reviewers. Since the rise of the research laboratory in the late nineteenth century, enthusiastic laboratory scientists, drug manufacturers, and doctors had proposed numerous medical interventions, some of which worked, and some of which didn't work or caused harm. In the spirit of Archibald Cochrane, EBM proposed a particular solution to the problem: the generation and compiling of therapeutic evidence and the "critical appraisal" of the evidence. Thus, while laboratory medicine was often filled with therapeutic optimism, contemporary epidemiological medicine exhibits, as a counterreaction, a kind of therapeutic skepticism, in that it is generally unconvinced by the promise exhibited by interventions in the laboratory until they have passed the test of high-quality, unbiased epidemiological research.

In the early twenty-first century, a related strain of skepticism about contemporary treatments has been cultured by medical insiders and outsiders who are concerned about the reliability of medical evidence due to bias seeded in it by its most common sponsor, the pharmaceutical industry. These worries are not new but have grown alongside the increasing role of the drug industry in clinical trials research since the middle of the twentieth century (Bothwell et al. 2016). Several critics sounded the alarm about pharmaceutical industry bias.[2] Evidence of industry bias comes partly in the form of investigations into its mechanisms, such as publication bias (the preferential publication of studies with favorable results), and partly from meta-research findings of ominous trends in the medical literature, including a positive association between flattering therapeutic study outcomes or conclusions and funding from companies producing the therapies being flattered (Lundh et al. 2017; Hansen et al. 2019). While various proposals seek to prevent biases like these in future research, it remains unclear how to account for their influence on our current medical evidence during evidence evaluation.

In answer to the latter problem, Jacob Stegenga (2018) defended a contemporary therapeutic skepticism that he calls Medical Nihilism. Medical Nihilism is the thesis that we should have low confidence in the effectiveness of medical interventions today. It is supposed to follow from the widespread existence of bias, among other problems in medical research. In this chapter, I will develop a distinct therapeutic skepticism that I call Medical

[2] For example, Sheldon Krimsky (2003), Marcia Angell (2005), Ben Goldacre (2012), David Healy (2012), and Peter Gøtzsche (2013).

Meta-Evidentialism as an alternative response to troubling concerns about biased or misleading therapeutic evidence and as a thesis about how these concerns should influence our beliefs about the effectiveness of medical interventions. Meta-Evidentialism is the view that we should be less optimistic about the magnitude of treatment effects measured in therapeutic research, lowering our confidence in the effect estimate, after learning of meta-evidence such as meta-research evidence of bias in the medical literature. Given that therapy—and not just theory—is a central concern in medicine (Broadbent 2019), the question of what position to take on the accuracy of therapeutic evidence and the effectiveness of medical interventions should be seen as fundamental in the philosophy of medicine, akin to the problem of whether scientific theories are accurate in the philosophy of science (Fuller 2021a).

Rather than Medical Nihilism, Medical Meta-Evidentialism is the most reasonable response to the problem of biased medical evidence. In section 8.2, I introduce meta-research findings of bias in the medical literature, the concepts of biased and misleading evidence, and Medical Nihilism. In section 8.3, I clarify Medical Nihilism and advance three objections: the broad-brush objection, the multiple measures of effectiveness objection, and the vague-probabilities objection. In section 8.4, I propose my alternative, Medical Meta-Evidentialism, which avoids the objections to Medical Nihilism. In section 8.5, I defend Meta-Evidentialism via an argument from meta-research evidence.

## 8.2. Biased Evidence and Medical Nihilism

Philosophers and scientists view several features of our current medical research landscape, including pharmaceutical industry sponsorship of research, as a major epistemic threat to medical evidence. While reactions to the problem of biased therapeutic evidence are varied, one strong response is Medical Nihilism. In this section, I discuss the problem of biased or misleading evidence and lay out Stegenga's main argument for Medical Nihilism.

A recurring theme in the EBM literature is the concern that clinical research is biased, owing to methodologic problems as well as contextual problems such as industry sponsorship. One classic (and iconoclastic) paper about the methodologic and social threats to medical evidence is John Ioannidis's (2005) "Why Most Published Research Findings Are False."

Ioannidis models the influence of bias, small studies, small apparent effect sizes, unselective hypothesis testing, research flexibility, and financial interests on the probability of research findings being veridical and concludes that most published research findings, including findings from the biomedical sciences, are not.[3] That these problems are not merely theoretical or uncommon aberrations in medical research is supported by findings from the discipline of meta-research. Meta-research is a field that studies trends in research, including the pervasiveness of potential biases such as methodological flaws, publication bias, conflicting study results, and the effects of industry sponsorship. Ioannidis, who is perhaps its figurehead, describes the field as "conducting scientific research on research itself" (Ioannidis et al. 2015, 2). Meta-research relies heavily on epidemiological methods to quantify and analyze research trends, including systematic review and meta-analysis.

Here are some notable examples of the kinds of results that the field is delivering. A random sample of over 1,500 interventions included in the Cochrane Database of Systematic Reviews found that among first-listed primary outcomes, only 10% are supported by high-quality evidence according to the popular GRADE approach to rating evidence (Howick et al. 2022). In the literature on publication bias, a meta-analysis found that trials with positive results are 1.78 times as likely to be published compared to trials with negative results (Hopewell et al. 2009). Meanwhile, a meta-study of systematic reviews of RCTs investigating healthcare interventions found that 38% change at least one outcome from the protocol to the published review. Such changes can result in bias if they are not methodologically indicated, but they were rarely explicitly justified (Page et al. 2014). A meta-analysis of industry sponsorship effects found that industry-sponsored drug and device studies are 1.27 times as likely to report favorable efficacy results and 1.34 times as likely to report favorable conclusions compared with non-industry-sponsored studies (Lundh et al. 2017). Finally, a study of systematic reviews found that reviews with financial conflicts of interest are twice as likely to come to favorable conclusions (Hansen et al. 2019).

Sergio Sismondo (2008) documents mechanisms of pharmaceutical company bias, including study design biases, marketing trials, scientific misconduct, favorable interpretations of trial data, publication bias, and the ghost management by drug companies of the entire research process, from

[3] For a critical response to Ioannidis's paper, see S. Goodman and Greenland (2007).

planning to publication. Likewise, Ben Goldacre (2012) investigates endless drug company tactics under the headings of ineffective drug regulation, biased trials, and drug marketing. But his most in-depth research examines the pervasive phenomenon of missing trial data. In one of his most unsettling anecdotes, reboxetine, an antidepressant, was compared to placebo in seven clinical trials. The only trial subsequently published in an academic journal was the one study that had a positive result; the other six studies involving almost ten times as many patients showed negative results and went unpublished (Eyding et al. 2010).

These kinds of amply documented research problems can result in therapeutic evidence that is misleading, potentially influencing evidence users to make therapeutic inferences that are unsound. There is an important connection between the concepts of "misleading evidence" and "biased evidence." According to the 2011 *Cochrane Handbook* for systematic reviewers, research bias is "a systematic error, or deviation from the truth, in results or inferences" (Higgins et al. 2011).[4] (Bias is defined as "systematic" to distinguish it from random error, another common source of deviation from the truth.) On this definition, inferences from medical evidence that deviate from the truth—or are unsound—can be considered to suffer from bias. In other words, bias is a feature both of misleading evidence and of inferences made from misleading evidence.[5]

The conceptual connection between biased medical evidence and biased medical inference forged by this definition dovetails with the argument theory or inference theory of medical evidence from Chapter 5. On the argument theory, medical evidence establishes a premise in an argument or chain of reasoning for a medical hypothesis. For instance, data from a randomized trial showing a difference in outcome between two study groups establishes an empirical premise in a causal inference about the effect size in the trial population. Say that the study suffers from selection bias: there is a baseline difference in confounders between study groups that explains the difference in study outcome. The difference in study outcome, as evidence for a premise

[4] More recent versions of the *Cochrane Handbook* exclude the word "inferences" from this definition (Boutron et al. 2024). The authors also argue that "[b]ias should also not be confused with the external validity of a study." This revised definition does not allow us to speak of "biased predictions" nor of "biased causal inferences" but only of biased study results. It is thus most consistent with the measurement theory of evidence introduced in Chapter 5, in which medical evidence consists of more biased and less biased measurements of clinical outcomes.

[5] For other accounts of the nature of bias or research bias, see Thomas Kelly (2023), Reutlinger (2020), and Wilholt (2009).

in the causal inference, is misleading evidence because the causal inference is unsound, and both the misleading evidence and the causal inference made from it are thereby biased.[6]

One could even go an inference further. What makes the misleading trial results misleading *medical* evidence on the argument theory is that they support a premise in a sub-argument that is part of a more complicated chain of reasoning warranting a *medical* hypothesis, in this case, a hypothesis about a patient's response to the intervention. Thus we could say that the other inferences that are part of this reasoning are also likely biased by the misleading evidence because they also will likely be unsound. For instance, in the Risk Generalization-Particularization Model of prediction introduced in Chapter 5, the conclusion about the effect size in the study population is a premise in a Generalization inference, which concludes that the same effect size would occur in some target population. If the causal inference is unsound, its causal conclusion is likely untrue, and thus the Generalization inference is likely also unsound because it starts from a false premise. On the argument theory, we can speak of the trial evidence as *indirect* misleading evidence for the conclusion about the target population, and thus we can say that this extrapolation inference is also biased by the evidence.

This capacious inferential conception of research bias also permits us to diagnose bias resulting from lack of study-target comparability (or lack of external validity) rather than from a confounded study association (or lack of internal validity). Say that the association between intervention and outcome found in the trial now genuinely reflects the effect of the intervention on the outcome. Our causal inference is sound and the conclusion about the effect size in the study population is true. However, if the study population is not sufficiently similar to the target population, then the Generalization is nonetheless unsound even though the causal inference is sound. Because the trial evidence indirectly supports this unsound inference, it is misleading and biases our extrapolation.

An inferential conception of bias draws attention to the various ways that study results can mislead those making medical inferences, including through "causal inference bias," "extrapolation bias," or both. However, the

[6] The causal inference for which RCT results provide evidence can be constructed as a deductive argument (see section 8.5). Nancy Cartwright (2007) argues that this is the advantage of the RCT method, which she thus calls a "clincher." However, comparative observational study results can support the same causal inference. The relative advantage provided by RCT results according to EBM proponents is that the results are more likely to be unbiased—that is, the causal inference is more likely to be sound.

inclusive concept has an important drawback. While biased causal inference results only from features of the study, biased extrapolation results partly from features of a target population, and while it is unambiguous to say that a causal inference is biased when we are talking about a single study population, it is ambiguous to speak of study results as resulting in biased extrapolation because extrapolations about some target populations may be sound while extrapolations about other target populations are unsound. To avoid this ambiguity, a defender of the inferential perspective on research bias could stipulate that study results are "biased" only when they lack external validity for common target populations. (For other target populations, the evidence may be misleading but is not "biased.") Or they could stipulate that only therapeutic evidence that misleads our *causal inference* is biased. (For other inferences, evidence may be misleading but is not "biased.") Or misleading research results themselves are not "biased," only unsound inferences made from them are. Or finally, just as data are evidence only relative to a particular inference (per the inference theory of medical evidence), evidence is biased only relative to a particular inference.

Whichever stipulation we choose, we should follow Daniel Steel (2018) in accepting that the relevant inferences for diagnosing bias include those inferences made by the audience for the research and not only by the study authors. After all, it is common practice to consider the results of a comparative intervention study biased if, for example, there are important baseline differences between study groups that explain the difference in outcome—even if the study authors acknowledge this point in the results or discussion sections of the research report and do not themselves make false inferences from the results. The results are still biased because it is common to infer that a difference in outcome measured in a comparative intervention study is an effect of the intervention. The reader of the research could therefore be misled by the results (which is why it is important to call out bias in a research report whenever it is suspected). Evidence is biased when it is liable to support an unsound argument, regardless of whether anyone actually makes the inference represented in that argument.

Commonly, research bias is associated with a particular methodological feature such as study group allocation resulting in assortment of study participants into groups in a way that is correlated with the study outcome ("allocation bias"). Moreover, research bias is most associated with threats to the internal validity of a study (as in the example of allocation bias). On a narrow understanding of research bias limited to internal validity and

identifiable methodological flaws, certain meta-research findings might reveal that therapeutic evidence is misleading, but not necessarily due to bias. For example, as suggested above, "industry bias" can be due to a diverse set of mechanisms. These mechanisms can include the evaluation of treatments in populations or against comparators that are unlike the target populations or treatment comparisons that clinicians face in routine practice, thus threatening the *external* validity of therapeutic study results.

On the other hand, EBM sources sometimes do allow a wide range of problems, including "industry sponsorship bias," to count as biases, even while acknowledging that they can result from multiple mechanisms and may involve problems such as unrepresentative study populations (Catalogue of Bias Collaboration et al. 2019). Some authors even refer to "risk of generalizability biases" or "generalizability bias" (Pressler and Kaizar, 2013; Beets et al. 2020). Furthermore, if we accept that evidence is biased whenever it is misleading for any potential evidence users, then because inferences about target populations and target patients are commonly made by evidence users from the results of an intervention study, it follows that systematic threats to external validity can also result in bias. Thus, all things considered, I will invoke a wider inferential concept of bias throughout this chapter that recognizes threats to external validity, including (potentially) industry bias.[7] Adrian Erasmus (2023) uses the term "meta-bias" to refer to meta-level biases such as industry bias and selective reporting or publication bias that can be realized by multiple lower-level methodological mechanisms, a term used similarly by Steven Goodman and Kay Dickersin (2011). However, the main problem that motivates this chapter is the problem of misleading evidence. Because "misleading evidence" is naturally understood broadly, not much will hang on how broadly "bias" is understood. Readers unhappy with my inclusive understanding of biased evidence can simply substitute the term "misleading evidence."

[7] Inmaculada de Melo-Martín (2019) argues against conceptualizing industry's influence in terms of bias, for several reasons. For one thing, industry's influence is putatively mediated by multiple specific mechanisms that each requires their own specific responses. However, I question whether industry bias is unique in being due to multiple mechanisms. For instance, selection bias—a paradigmatic research bias if there ever was one—can be realized by either of two general mechanisms: allocation bias (when experimenters decide participants' group allocation) and self-selection bias (when participants decide their own group allocation; Howick 2011). Moreover, there *are* proposals for combating industry bias that are agnostic to the specific mechanism of industry influence (Resnik and Elliott 2013; Holman and Elliott 2018), including those that shield the sponsor from study design, analysis, or reporting, which is meant to prevent the sponsor's influence regardless of whether that influence would be exerted through deleting data points or using comparator drugs at substandard doses (for instance).

Holman and Elliott (2018) review solutions proposed by philosophers to the problems of industry funding. Most proposals are designed to prevent detrimental industry influence rather than calculate its influence on study results after the fact (though Holman and Elliott note that statistical adjustment can be applied to a meta-analysis post hoc to correct for publication bias). A couple of exceptions are proposals by Resnik and Elliott (2013), Elliott (2014), and Solomon (2020a), which I will note in section 8.4. However, the most radical and well-developed philosophical response to the problem of misleading or biased medical research is Jacob Stegenga's Medical Nihilism: in his words, "We should have little confidence in the effectiveness of medical interventions" (2018, 167).

Despite what its name might suggest, Medical Nihilism is not about all of *medicine* but about contemporary medical treatments; it is a *therapeutic* nihilism. And Stegenga's is "not the thesis that there are no effective medical interventions" (2018, 184) but rather that we should have low confidence in the effectiveness of medical treatments generally. This lack of confidence permits exceptions. Stegenga does not want to deny that there are some truly effective medical interventions; such interventions, like insulin and antibiotics, Stegenga celebrates and refers to as "magic bullets," borrowing an expression coined by the physician-scientist Paul Ehrlich (1908). Stegenga's thesis is about our stock of medical interventions; it is not about this or that intervention, nor certainly about every particular medical intervention. I will further discuss how best to interpret Stegenga's thesis in the next section.

In arguing for Medical Nihilism, Stegenga draws on a range of considerations. The pathophysiological basis of most contemporary diseases is complex, and thus we should not expect most of our interventions to approximate magic bullets that fly right to the heart of disease and target it with high effectiveness and specificity. The medical research methods favored by EBM, especially randomized trials, meta-analyses, and evidence evaluation schemes, are malleable or flexible, and thus can be bent to the will of researchers. Industry sponsorship influences this bending in the direction of showing a positive study result, even when the intervention studied is ineffective. Finally, there is an abundance of evidence in favor of Medical Nihilism: a long history of reversals of judgment about effectiveness in the medical literature, evidence of industry bias, evidence of discordant or conflicting results of studies of the same intervention, and a plethora of examples in which an intervention is found to have a small apparent effect size in a

clinical trial. Stegenga argues that, taken together, the existence of these problems supports Medical Nihilism.

To bring together these various considerations and show their relevance for his thesis, Stegenga lays out the Master Argument for Medical Nihilism. The Master Argument uses Bayes's Rule:

$$p(H \mid E) = \frac{p(E \mid H) \times p(H)}{p(E)}$$

H is a hypothesis claiming that a medical intervention is effective. E is evidence relevant to H: specifically, the results of a clinical trial or meta-analysis studying the intervention. The conclusion of the Master Argument is that the probability of H given E, p(H|E), is generally low—the thesis Stegenga calls Medical Nihilism. Stegenga argues that the prior probability of H (p(H)) is "low"; the probability of E (p(E)) is "high"; and the probability of E given H (p(E|H)) is "low"; therefore, by Bayes's Rule, p(H|E) is "low." The conclusion that p(H|E) is "low" is meant to follow deductively given the values Stegenga assigns to the terms on the righthand side of the equation. If Stegenga is right, then to convince the reader that his argument is sound, all he must do is justify the values he assigns to the terms on the righthand side ("low" and "high"). The bulk of Stegenga's efforts are directed toward that task.

Stegenga argues that p(H) is low because we should have low confidence that an intervention is effective before encountering clinical research evidence E. His first argument relies on the complex pathophysiology of most diseases, which defies precise manipulation. His second argument relies on the phenomenon of medical reversal: many interventions historically thought to be effective were later rejected as ineffective; thus there is a reasonable probability that a given new intervention would be found ineffective on future testing.[8]

Next, p(E) is supposed to be high because we should expect research evidence to show positive results, even more so when that evidence is sponsored by a pharmaceutical company, and even when the intervention is not truly effective. Despite their vaunted objectivity, clinical research methods are malleable, and thus sponsors or investigators can fine-tune the methods to produce positive results regardless of whether the intervention works. Hence

[8] This argument is basically a "pessimistic induction" of the kind classically advanced against scientific realism in the philosophy of science (Laudan 1981).

the probability of favorable-looking evidence E is high due to the ubiquity of bias that favors the intervention.

Finally, p(E|H) is low because most apparent effect sizes measured in trials are small and because studies are frequently discordant, contradicting each other. Stegenga argues that if most interventions were truly effective, these phenomena would be uncommon, and thus the probability of finding them—assuming that the intervention is effective—is low. On the one hand, if an intervention is effective, then evidence would probably not show that it is relatively ineffective. (I will analyze this argument further in the next section.) On the other hand, if an intervention is effective, then the evidence would probably be concordant rather than contradictory. (This argument assumes that the studies are adequately statistically powered to detect effectiveness and that biases producing false negative results are uncommon.)

Stegenga thus draws the general conclusion that we should have low confidence or credence in the effectiveness of medical interventions, even after considering evidence relevant to their assessment—evidence that may appear to demonstrate that our interventions are effective but is misleading. I agree that our evidence is commonly misleading, but I disagree that Medical Nihilism is the right conclusion to draw.

## 8.3. Against Medical Nihilism

Stegenga canvasses several considerations in support of Medical Nihilism. While I agree that they are relevant to assessing the effectiveness of medical interventions (and will make room for them in my alternative, Medical Meta-Evidentialism), in this section I will discuss three objections to Stegenga's position.

Before launching into the objections, it will help to get clearer on H, E, and Medical Nihilism. Donald Gillies criticizes Stegenga's thesis because Gillies (2019, 273–274) reads it as claiming we should have little confidence in the effectiveness of medical interventions except for an open-ended list of exceptions that Stegenga does not delimit with a demarcation criterion: "Consequently, it looks as if medical nihilism has degenerated into the empty claim that we should have little confidence in the effectiveness of medical interventions, except for those medical interventions in which we should have more than a little confidence." However, Stegenga's arguments *do* imply

some guidance on when we should have higher confidence in a particular medical intervention: when it approximates the magic bullet ideal, when research on that intervention is not funded by industry, when its apparent effect size is substantial, and so on.

More important, Gillies's criticism seems to rest on a misreading of Stegenga's thesis: that Medical Nihilism is a claim about individual medical interventions that quantifies overall medical interventions. This interpretation reads Medical Nihilism as saying that we should have low confidence in all medical interventions (with exceptions). Instead, Stegenga's Master Argument is about a generic hypothesis H, "a hypothesis that a medical intervention is effective" (J. Stegenga 2018, 176), and a better reading of Medical Nihilism is that it is about medical interventions as a group: "on average we ought to have a low posterior probability in that hypothesis, p(H|E)" (168). Medical Nihilism is about what our *average* credence should be across all medical interventions, or, derivatively, the credence we should assign to a generic intervention chosen from this group, assuming we know nothing else other than that it is a medical intervention. What exactly Stegenga means by "effectiveness" is a question I will address in section 8.3.2.

What about the content of E? Stegenga uses the preponderance of small effect sizes of medical interventions to argue that p(E|H) is low. Thus E is best understood as the finding of a small apparent effect size, a positive association between intervention and outcome of small magnitude in an RCT or meta-analysis of RCTs. Characterizing E in this manner allows us to relax a constraint imposed by Stegenga that p(E|H) must exceed p(E). Stegenga imposes this constraint because he stipulates that E provides some (weak) confirmation to H, and thus p(E|H)/p(E) must be greater than 1 so that p(H|E) is greater than p(H) by Bayes's Rule. However, if we understand E as the finding of a small association rather than as weakly confirmatory evidence, then we can drop this constraint because bias could be a better explanation for a small association rather than effectiveness, and thus a small association might not be confirmatory.

However, as Alex Broadbent (2019) argues, Stegenga's argument for a low p(E|H) also implies that H cannot be the hypothesis that the medical intervention has just any degree of effectiveness, or else H would render the probability of E (evidence consistent with a small effect size) high rather than low. Rather, H must assert that the intervention has an effect size greater than "small."

Having clarified Medical Nihilism and the Master Argument, I will now present three objections: a broad-brush objection, a multiple measures of effectiveness objection, and a vague-probabilities objection.

### 8.3.1 Broad-Brush Objection

My first objection is directed at the thesis Medical Nihilism itself. The *broad-brush objection* is the worry that Medical Nihilism is an overly general position. Something like this objection is perhaps the most common criticism of Medical Nihilism.[9] For instance, Broadbent (2019) raises a "blue bus objection" against Medical Nihilism: we should not apply general evidence to an individual medical case (analogous to the problem of applying general statistical evidence to legal cases in the courtroom). Similarly, Mathew Mercuri (2018) argues that we should not lower our confidence in a particular intervention just because the problems Stegenga identifies can or often do occur. Finally, Miriam Solomon (2020b) notes that certain areas of medicine, such as pediatric oncology, are home to effective therapy, while Arjun Devanesan (2020) provides other examples. One could multiply examples of effective medical interventions: many cancer treatments used in adult medicine, including various surgeries and chemotherapies; surgeries to remove inflamed gallbladders and appendixes, to extract advanced cataracts, and to replace failing body parts; thrombolytics for heart attack and ischemic stroke; anti-inflammatory and immunomodulatory agents with various indications, such as steroids and biologics; many drugs for symptom control, such as antiepileptics and heartburn medications; many reproductive technologies, such as birth control and fertility treatments; many vaccines and antimicrobials for viral, bacterial, and parasitic infections; and various life-sustaining interventions administered in emergency departments or intensive care units, from the mundane (such as intravenous fluids) to the heroic (such as mechanical ventilation and ECMO). In short, we should not paint medicine black with such a broad brush.

[9] See Broadbent (2019), Mercuri (2018), Devanesan (2020), Solomon (2020b). In contrast, Upshur and Goldenberg (2020) argue that Medical Nihilism rests on a misunderstanding of medical evidence. Broadbent (2019) also objects that while Medical Nihilism is advertised in its name as a nihilism about medicine, it is really only a therapeutic nihilism, and therapy is not, Broadbent argues, the core business of medicine.

There are two versions of this objection, one sturdier than the other. The version stated above is normative. However, a different version of the worry is defended by Devanesan when he writes that the heterogeneity of evidence in medicine renders a general conclusion like Medical Nihilism "highly implausible" (2020, 2); it is inaccurate to paint medicine black with such a broad brush. If Stegenga's thesis were a universal quantification over medical interventions (along with a well-delimited set of exceptions) claiming that we should have low confidence in each of them, then I would agree with Devanesan because evidence that is heterogeneous (in the sense that some evidence greatly supports some interventions, while other evidence does not greatly support other interventions) cannot warrant such a generalization. However, as I just argued, Stegenga's thesis is best construed as a claim about our confidence in the effectiveness of medical interventions averaged across interventions. Our average confidence *can* be low even when our confidence in some medical interventions is high.

As an illustration, Figure 8.1 shows two (of many) possible distributions for our confidence in the effectiveness of individual medical interventions. Distribution A is narrow and unimodal: our evidence generally supports a low confidence in a particular medical intervention. Distribution B is wide and double-humped: while most of the time our evidence supports a low

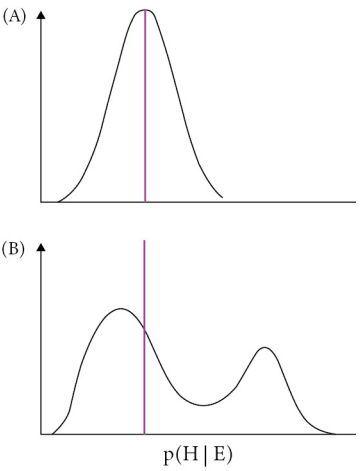



**Figure 8.1** Two possible scenarios, A and B, representing what our confidence should be in individual medical interventions. x-axis: our confidence represented by p(H|E), ranging from 0 to 1; y-axis: frequency of medical interventions; vertical line: average confidence.

confidence in a medical intervention, much of the time it supports a substantially greater confidence. Yet despite the heterogeneity, we can still plot an average for distribution B just as we can for distribution A.

A sturdier version of the broad-brush objection insists that we *should not* adopt an average confidence even if we can; we *should not* use such a broad brush. Medical Nihilism is a normative thesis: "[O]n average we ought to have a low posterior probability in that hypothesis, p(H|E)" (J. Stegenga 2018, 168); "[w]e should have little confidence in the effectiveness of medical interventions" (167). The broad-brush objection provides a normative response. While an average is a useful statistic if distribution A best represents our range of confidences in medical interventions, the average is not a very useful statistic if distribution B best represents our lot. The former situation is accurately described as therapeutic *nihil*ism: we should have low confidence in most medical interventions and moderate to high confidence in almost *nihil*. In contrast, we can call the latter situation therapeutic *mixed-bag*ism: we should have a range of confidences in medical interventions, varying from low to high (which could in principle average out to a low confidence).

Mixed-bagism is a more plausible representation for contemporary medicine than therapeutic nihilism. (Distributions that are more optimistic than the one in Figure 8.1B are also more plausible.) Stegenga agrees that there are many examples in medicine of magic bullets that are like insulin and penicillin: highly effective and specific (2018, 184). Yet Stegenga's main thesis is a nihilistic one, which would be more appropriate given distribution A rather than distribution B. Medical Nihilism paints over medicine's mixed success.

To make his thesis more discriminating and paint with a finer brush, Stegenga could advertise therapeutic mixed-bagism as his main conclusion rather than Medical Nihilism. Or he could narrow the scope of Medical Nihilism. Given that among treatments for infectious diseases there are many magic bullets—antibiotics for bacterial infections, antivirals against HIV and COVID-19, numerous effective vaccines—Stegenga could restrict his nihilism to noncommunicable disease treatments. Furthermore, given that his criticisms of medical research best apply to contemporary pharmaceutical research (as Broadbent 2019 notes), Stegenga could further restrict nihilism to contemporary drugs. These are potential avenues for crafting a generalization that captures the depressing left hump in distribution B. All of this is a sophisticated way of stating a plain point: given that our medical interventions are a mixed bag, we should not endorse any thesis like Medical

Nihilism that paints in broad strokes when a finer, more discriminating conclusion is warranted. In section 8.4, I will provide an alternative to Medical Nihilism that lets us better discriminate therapeutic winners from losers.

### 8.3.2 Multiple Measures of Effectiveness Objection

My next objection, the *multiple measures of effectiveness objection*, concerns Stegenga's argument from the preponderance of seemingly small effect sizes. In short, Stegenga draws a line between categorically effective and categorically ineffective interventions, which is sensitive to how we measure effectiveness. The measure that Stegenga favors is the absolute risk reduction (ARR) of a medical intervention. However, other measures would better quantify the kind of effectiveness that Stegenga is after (clinical value), and, using those measures, contemporary medical interventions might not look so ineffective.

Stegenga notes that effectiveness comes in degrees and takes the ARR to be a good metric for quantifying effectiveness. Recall that the ARR is the difference in absolute risk of the outcome between comparison groups. It can range from 0% to 100%. However, Medical Nihilism is a thesis about our confidence in the categorical effectiveness of medical interventions. In the Master Argument, H is the hypothesis that a medical intervention is (categorically) effective. Stegenga must draw a line between medical interventions that are effective and medical interventions that are ineffective. Recall that his argument from the ubiquity of small effect sizes requires that an intervention with a small effect size is not considered effective (so that a small apparent effect size E renders p(E|H) low). Thus Stegenga draws the line between effectiveness and ineffectiveness such that interventions with a small ARR are on the ineffective side. While he does not provide a numerical cut-off for effectiveness, he lists several examples of interventions with "numbers needed to treat (NNT)" in the range of 73–104, which is equivalent to an ARR of 1.0–1.4% (2018, 174).[10] However, I do not need to assume a particular numerical cut-off to advance my objection; I can grant for the sake of

[10] The NNT is the inverse of the ARR and is interpreted as the number of individuals one would have to treat with the intervention to prevent one outcome (in the long run). I did a survey of medical interventions listed on the website NNT.com (most of which are drugs) and found that of the 70 deemed to be effective and with listed ARRs, 35 (50%) had an ARR less than 10% and 53 (76%) had an ARR less than 20% (Fuller 2020b).

argument that many medical interventions today have small ARRs (however we want to define "small").

The ARR or the NNT (which is simply the inverse of the ARR) is just one way of quantifying an intervention's effectiveness, what is known as an "absolute outcome measure." Another set of outcome measures one could use are "relative outcome measures," such as the relative risk (RR) or the relative risk reduction (RRR). The ARR and the RRR sometimes have vastly different magnitudes for the same intervention, outcome, and population. (Recall the example of statins for heart attack prevention in a low-risk population from Chapter 6, where the ARR was 1%, while the RRR was 33%.) This illustrates that Stegenga's argument from the ubiquity of low effect sizes depends on his choice of measure for quantifying effectiveness. With Jan Sprenger, Stegenga provides three arguments that justify the use of absolute outcome measures over relative measures: relative measures are often confused with absolute measures, leading to overestimation of effectiveness (meaning: effectiveness quantified using an absolute measure); absolute measures satisfy several desiderata for a measure of causal strength that relative measures fail to satisfy; and absolute measures are more informative than relative measures for decision-making, according to expected utility theory (Sprenger and Stegenga 2017). I will examine the implications of this third justification for Stegenga's small effect size argument.

The expected utility justification for preferring absolute over relative measures demonstrates that, given information on the relative utilities (the magnitude of value and disvalue attached to various outcomes), the ARR or NNT allows a decision-maker to determine whether the expected utility (EU) of one intervention (the probability-weighted sum of the utilities) exceeds the expected utility of another intervention or no intervention. In their words, the ARR and NNT are "EU-sufficient." Meanwhile, the RRR and RR are not EU-sufficient, and so are less informative regarding the expected value of a medical intervention. However, there are better ways of quantifying the value of a medical intervention than using an intervention's ARR alone. In the spirit of Stegenga's endorsement of the ARR based on its EU-sufficiency, we should prefer these other metrics to the ARR. And they will sometimes (perhaps often) cast the medical intervention in a better light.

First, interventions that have what we might consider to be a small ARR can have a substantially greater expected utility given the importance of the outcomes they prevent. According to the reasoning Stegenga uses to justify using the ARR over the RRR, invoking the former's greater information content for expected utility calculations, we should prefer the difference

between the expected utilities, the expected utility difference (EUD), to the ARR because it allows us to compare interventions given even less information; it is even more informative (perhaps "EU-super-sufficient"). The EUD is thus a better tool than the ARR for measuring the value of a medical intervention, and from the perspective of the EUD, many of our medical interventions might come out looking more effective due to the great disutility of the outcomes they prevent. For example, statins, antihypertensives, and anticoagulants/antiplatelet agents prevent heart attacks and strokes that lead to death and disability. We should not consider these interventions to be categorically ineffective on the basis that they have a small ARR (whatever other reservations we might harbor about these interventions).

Moreover, interventions with a small ARR can contribute to important cumulative gains in three respects. First, preventive medical interventions are often used not in isolation but as part of a suite of interventions, particularly in secondary prevention. For example, multiple interventions are typically used in the acute heart attack setting, in the secondary prevention of further heart attacks or strokes after an initial event, and in multimodal cancer therapy. In isolation, each intervention might have a smaller effect, but their collective effect is greater. For example, recall from Chapter 6 that in a certain population with invasive breast cancer, the ARR for death with trastuzumab was 5% at fifteen years. However, this 5% isolates the effect of trastuzumab, which is usually added to a regimen of other standard interventions. In the same population, the ARR from the combined multimodal treatment regimen is predicted to be 30%, which excludes the benefit of the surgery to which the regimen is added (Figure 8.2).

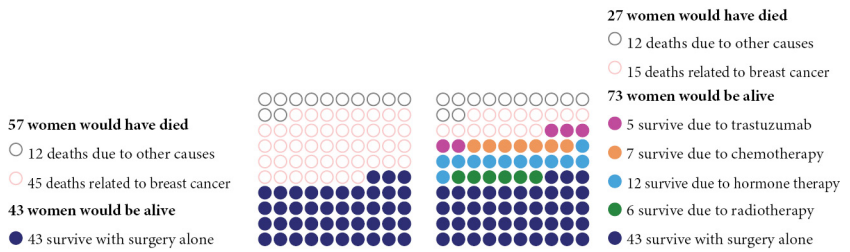



**Figure 8.2** Decision aid representing deaths in certain women with invasive breast cancer at fifteen years following breast cancer surgery. The two displays compare the population treated with a standard multimodal treatment regimen (right) to the population not treated with the regimen (left).
National Health Service (2024).

Second, many interventions in areas like pediatric oncology represent incremental gains in effectiveness over the standard of care to which they are compared (Solomon 2020b), for instance, a newer chemotherapy protocol compared to an older one. Newer chemotherapy might have a small effect as measured through comparison with older chemotherapy ("comparative effectiveness"), but these incremental gains can cumulate and lead to substantial benefit of the newer chemotherapy compared to no chemotherapy.

Finally, preventive interventions with a small ARR are often widely indicated. In fact, it is partly because they are used in wide populations having a low baseline risk of an adverse event that their ARR is small (the ARR can be calculated by multiplying the RRR by *the baseline absolute risk*). They can nonetheless have a large cumulative impact at a population level, preventing a large absolute number of outcomes, because this low-risk population is often large (and the absolute number of outcomes prevented is calculated by multiplying the ARR by *the size of the population*).[11] As an illustration, one study modeling the impact of COVID-19 vaccination campaigns worldwide over 2021 estimated an ARR of death from COVID-19 among people in 185 countries and territories of approximately 0.2281% or 0.3121%, depending on how deaths were estimated. Because the world is a vastly populated place, these figures represent 14.4 million or 19.8 million deaths averted, respectively (Watson et al. 2022). All three kinds of cumulative gains discussed above can translate into marked improvement in measures of population benefit such as life expectancy.[12]

[11] Stegenga (2018, 124) anticipates this objection: "For example, if a cholesterol-lowering drug has a 1 percent absolute reduction in the risk of death, and ten million people consume the drug, then 100,000 lives are saved. That is a great outcome." He responds by adopting the individual perspective: "However, it is not obviously great from the perspective of a particular typical patient. . . . An individual patient and her physician want to know that if they use a particular intervention then there is a reasonably good chance that the intervention will be effective for this particular patient". This response leaves open the possibility that one could accept Medical Nihilism and yet maintain that medicine provides great population-level benefit. The epidemiologist Geoffrey Rose (1992) referred to the phenomenon that an intervention could bring great benefit to a large population (in terms of the number of cases or deaths averted) while bringing little benefit on average to each of the population's individuals as the "prevention paradox."

[12] One analysis estimated a total life expectancy gain of five years in the U.S. during the 1900s attributable to widely used medical interventions, with three of those years of medicine-attributable increased life expectancy occurring after 1950 (Bunker et al. 1994; Bunker 1995). These statistics do not capture gains in function and quality of life from a diverse range of medical interventions, including many of the examples listed in section 8.3.1. On the reverse side, after around 2000 increasing opioid addiction began to impair quality of life for many and even exert a downward force on life expectancy in the U.S., spurred partly by the mass marketing of OxyContin, a medical intervention (Quinones 2015; Case and Deaton 2020).

When evaluating the clinical and public health impact of medical care, we should consider the disutility of disease outcomes being treated as well as the cumulative impact of our interventions. Granted, I have not conclusively shown that were we to do so, medical interventions would on average come out looking impressive. While we have easy access to estimates for the ARR, we generally do not have easy estimates for the EUD, which requires estimates for a range of utilities. Moreover, there are important costs, side effects, and harms of medical interventions to weigh in calculating the EUD, which are often underestimated and underreported (J. Stegenga 2018; Howick et al. 2022). Finally, the harms of these interventions will cumulate in similar ways to their benefits. In fact, multi-intervention regimens will sometimes result in harmful treatment-treatment interactions that don't appear when these interventions are used in isolation. (On the other side, beneficial treatment-treatment interactions can also occur.) The primary ambition of the multiple measures of effectiveness objection is to point out that there are more comprehensive ways of quantifying the value of medical care that would *plausibly* register more substantial "effectiveness," which should cast some doubt on Stegenga's premise that medical interventions generally have a small or insignificant effect.

To summarize, Stegenga draws a line between categorically effective and categorically ineffective interventions using the ARR and considers interventions with a "small" ARR to be ineffective. However, other measures of clinical value should be used in lieu of or in addition to the ARR of individual interventions (such as the EUD or the cumulative impact of multiple interventions in reducing the absolute number of disease outcomes), and they would plausibly register a larger benefit of treatment. This would weaken Stegenga's argument from a small effect size crisis: because the magnitude of apparent benefit captured by evidence E in his argument is now less often "small," on average the $p(E|H)$ is higher. In section 8.4, my alternative to Medical Nihilism avoids these difficulties because it is less concerned with categorical effectiveness and less committed to particular measures of effectiveness.

### 8.3.3 Vague-Probabilities Objection

My final objection to Medical Nihilism concerns the overall Master Argument. The *vague-probabilities objection* argues that Stegenga assigns

values to the probabilities in the Master Argument that are problematically vague, yet it is difficult to make them more definite. Unfortunately, the use of vague probabilities makes the Master Argument difficult to interpret and evaluate.

Medical Nihilism itself uses a vague term to describe our confidence in the effectiveness of medical interventions: "low" or "little." There is no regimented meaning of terms like "low" or "high" in medicine, which makes the implications of Medical Nihilism difficult to discern.[13] Doctors and patients might desire to use a treatment in which they have low confidence if the expected benefit/harm profile is better than the alternatives (as they should, according to the same decision theory that Stegenga uses to defend the use of absolute outcome measures).

However, the main thrust of my vague-probabilities objection concerns the other probabilities in the Master Argument, the terms on the righthand side of Bayes's Rule. Stegenga assigns the values "low" or "high" to these terms. But given the vagueness of expressions like "low" and "high" in a context with no established conventional meaning, it is difficult to interpret their meaning and to evaluate Stegenga's arguments for the values he assigns. For instance, is the p(E) so high that virtually no finding of a positive outcome in a clinical trial or meta-analysis can confirm the hypothesis, or only high enough that most evidence is confirming but just not quite as confirming as commonly thought? This kind of distinction is important when applying Bayes's Rule. Saying that the p(E) is "high" does not discriminate between these possibilities. Stegenga's argument that there are ubiquitous biases in medical research that favor finding a positive outcome does not answer this question either.

Stegenga writes that "there can be no precise or general answer" to how low is "low" (2018, 183), and "I do not pretend that the terms of the master argument can be precisely quantified" (178). However, the problem is not that "low" is imprecise but rather that it is vague. By "vague," I mean that no range of definite real numbers is associated with the expression, in comparison to imprecise values, which *are* associated with a (wide) range or interval of definite real numbers. (This is the sense of imprecision represented by the confidence interval in clinical epidemiology.) While we can insert the upper and lower bounds of probability intervals into the terms on the righthand side of

[13] The dangers of vague probabilistic expressions led the U.S. Intelligence Community to endorse using numbers to express probabilities in intelligence communications, even for hypotheses that concern unique events (Tetlock and Gardner, 2016).

Bayes's Rule to calculate a posterior probability range,[14] we cannot calculate anything from vague linguistic expressions. If Stegenga's Master Argument was imprecise, its conclusion would be imprecise, but at least it might follow. Because it is instead vague, no clear conclusion follows. Stegenga could clear up the Master Argument's vagueness by replacing "high" and "low" with numerical ranges. But then two problems present themselves: What constraints should govern the choice of numerical ranges, and could Stegenga defend certain numerical ranges given those constraints?

With regard to the first question, Stegenga suggests that the Master Argument is deductively valid because it uses Bayes's Rule. Because the argument is meant to be deductive, ensuring its validity is essential. However, we can guarantee that the Master Argument is valid only given certain constraints on "high" and "low." Here's an illustration of the underlying issue. If the range for a "low" probability is 0–0.5, while the range for a "high" probability is 0.25–1.0, then when we plug in the "low" range for p(E|H) and p(H), and the "high" range for p(E), the posterior probability range is 0–1.0. This posterior range includes values in both the "low" and "high" range. Therefore the posterior probability might not be "low," even when the p(E|H) and p(H) are "low" and the p(E) is "high." Following Stegenga's strategy of making the Master Argument deductively valid, we can impose a validity constraint on the values for "high" and "low": the range for "low" is entirely below the range for "high." Then, given a "low" p(E|H), a "low" p(H), and a "high" p(E), the p(H|E) is always "low": the upper bound of the posterior range is at least as low as the upper bound of the range for p(E|H) and p(H).[15]

Adopting this constraint, the Master Argument can be defended by arguing for some range for p(H) and p(E|H) and some range for p(E) such

[14] I read Stegenga as using terms like "low" and "high" to stand in for a range of reasonable real numbers to which precise Bayesians can help themselves in choosing a single precise probability, rather than as standing in for an imprecise probability (a real number interval) that an imprecise Bayesian should adopt. However, nothing hangs on which reading we choose because vagueness is a problem for both precise and imprecise Bayesians insofar as both rely on conditionalization using mathematical functions. To be clear, I'm not denying that our relevant credences are vague or that they ought to be vague. But if we are to use Bayes's Theorem, we must at least *model* our credences using precise probabilities, imprecise probability intervals, or some quantitative alternative (Lyon 2017).

[15] If we take "low" and "high" to fully exhaust the probability space, then "low" is [0,x] and "high" is (x,1]. Plugging in the "low" range for p(E|H) and p(H) and the "high" range for p(E), the range for p(H|E) is then [0,x], which is "low." The validity constraint assumes that "low" and "high" are being used consistently throughout the Master Argument, or at least that "low" is always less than "high." If we loosen this assumption, however, and allow that "low" values can be greater than "high" values within the very same argument, then "low" and "high" lose all meaning.

that the former range is lower than the latter. Unfortunately, the prospects of successfully mounting such a defense are not good. Doing so requires quantitative evidence that allows us to assess the relative values of these terms. Stegenga marshals empirical evidence in support of his Master Argument, including the kinds of meta-research findings I discussed in section 8.2. I have argued that these kinds of evidence are relevant for assessing the effectiveness of medical interventions (Fuller 2018a). However, I have been unable to successfully use them to constrain the values of the terms of the Master Argument enough to defend the argument (Fuller 2020b).[16] That is because these quantitative findings do not allow us to directly estimate (precise or imprecise) values for the terms on the righthand side of Bayes's Rule, or even their relative magnitudes.

For instance, the finding discussed in section 8.2 that industry-sponsored drug and device studies are 1.27 times as likely to report favorable efficacy results (Lundh et al. 2017) suggests that industry-sponsored evidence is prone to bias that can generate a positive result E. This finding thus supports a p(E) greater than zero because many studies are industry-sponsored and some of them will have a positive result regardless of whether the intervention studied is effective. But the industry bias finding doesn't allow us to narrow in on a lower limit for the p(E), which is what we need to determine whether p(E) is greater than the terms in the numerator. Like most of the evidence and considerations that Stegenga uses to defend Medical Nihilism, the meta-research evidence has the effect that learning of it should often lead us to alter our credences (in this case, raising our p(E) and lowering our p(H|E)), but it does not support specific (even quite imprecise) values for these credences. Yet for the Master Argument to be a convincing argument, the values assigned to its probabilistic premises must be justified over other values these terms could take, and that might entail a value for its probabilistic conclusion that is higher than the range we've defined as "low."

In summary, one cannot do Bayesian conditionalization with vague probabilities, yet the Master Argument attempts to do so. Stegenga could instead assign imprecise numerical ranges that obey a constraint on the deductive validity of the argument. But the soundness of such an argument

[16] Meta-research studies that are potentially relevant for assessing the reliability of therapeutic studies or are otherwise potentially relevant for the Master Argument include surveys of "medical reversals," or RCTs finding that standard medical practices are less effective than previous practices (e.g., Prasad et al. 2013; Herrera-Perez et al. 2019), as well as surveys of disagreements among the results of therapeutic studies (e.g., Ioannidis 2005; Pereira et al. 2012; Toews et al. 2024).

is difficult to defend, even using relevant empirical findings from the field of meta-research.[17] As we will now see, my alternative is motivated by the same considerations as the Master Argument, including the same empirical findings, but avoids the intractable difficulty of defending an argument about medical interventions taken as a group, instead seeking to evaluate interventions individually.

## 8.4. Medical Meta-Evidentialism

The problem of biased or misleading therapeutic evidence discussed in section 8.2 demands that we adopt a critical attitude toward therapeutic study results. Medical Nihilism provides one solution. I propose an alternative brand of therapeutic skepticism, Medical Meta-Evidentialism, that avoids the problems with Medical Nihilism explored in the previous section.

Abundant empirical studies, including findings from the field of meta-research, suggest that the therapeutic evidence upon which doctors most commonly rely—from clinical trials and meta-analyses—is biased or misleading, partly because it often fails to live up to the highest methodological standards for research, but also because of social influences such as industry funding. Thus, in searching for the proper critical attitude toward therapeutic evidence, we cannot rely solely on EBM, which tells us when there is a risk of bias owing to the absence of certain methodological safeguards but not the actual influence of bias on study results. Moreover, EBM has historically offered fewer prescriptions and has been less effective in handling meta-biases, especially those such as industry bias that are due to aspects of the research context rather than a single specific methodological fault.

My proposed solution is *Medical Meta-Evidentialism: we should generally lower our confidence in an effect estimate from contemporary therapeutic research (and should generally raise our confidence in lesser effect estimates) after learning of meta-evidence about that research.*[18] By "contemporary

[17] See Fuller (2020b) for an earlier version of the vague probabilities objection, and J. Stegenga (2020) for a response.

[18] Medical Meta-Evidentialism can be rephrased as a thesis about the effect estimate rather than our confidence in the effect estimate: our confidence-adjusted estimate of the effect should generally differ from the unadjusted effect estimate provided by contemporary therapeutic research (and generally should be lower) after learning of meta-evidence about that research. This restatement assumes a kind of probability-weighting procedure that could be implemented in, for instance, a meta-analysis in which we assign a probability weight to the results of each included therapeutic study.

therapeutic research," I am referring to intervention studies such as clinical trials or meta-analyses of trials that are typically used to make first-order inferences about therapeutic effects (which, in the next section, I will call "first-order evidence"). I will explain further what I mean by "meta-evidence" in the next section, but I am referring roughly to second-order evidence about these intervention studies—including some of the results provided by the field of meta-research—that provides evidence of systemic bias in the medical literature. To put my thesis in the language of the argument theory of medical evidence, we should generally lower our confidence in a hypothesis about a particular magnitude of effect, a hypothesis that receives prior support from therapeutic evidence through an argument that warrants the hypothesis. Medical Meta-Evidentialism concerns hypotheses about the effect in the study as well as hypotheses about the effect in a target population because meta-evidence such as evidence of industry bias may be relevant to either or both settings.

Our effect estimate might simply give the net direction of the effect: "the effect is net positive." Put differently, "the intervention is effective." Thus Meta-Evidentialism implies we should generally lower our confidence in the effectiveness of medical interventions after learning of meta-evidence. It is a variety of therapeutic skepticism that leads us to doubt effect estimates from contemporary therapeutic research (and to have more confidence in less optimistic effect estimates). This skepticism is a challenge for the view that "high-quality" therapeutic studies generally provide an accurate measurement of the effect size and for evidence users who do not already modulate their confidence in those results using meta-evidence. When we add the bracketed component of the thesis, Meta-Evidentialism is also a skepticism about the interventions themselves, suggesting that we should generally believe that they are less effective than first-order evidence suggests. In comparison, Medical Nihilism, a different variety of therapeutic skepticism, claims we should have low confidence in the effectiveness of medical interventions rather than lower confidence.

I have bracketed part of Medical Meta-Evidentialism to draw attention to two different claims that the thesis makes, which require different justification. The unbracketed component is a commitment to conditioning our confidence on meta-evidence, which should generally change our confidence in a therapeutic hypothesis, just as learning about an additional piece of first-order therapeutic evidence should. My argument in the next section, the argument from meta-research evidence, is mainly an epistemic defense

of this unbracketed claim. However, on its own this commitment does not tell us whether we should reallocate the subtracted confidence to lesser effect estimates or to greater effect estimates; it doesn't tell us whether we should become more optimistic or less optimistic about our medical interventions. We can call the bracketed component of Meta-Evidentialism the "overestimation thesis." It adds to Meta-Evidentialism the commitment that generally meta-evidence reveals biases that lead to overestimating the effect. This is an empirical claim about meta-evidence, especially meta-evidence from the field of meta-research. Thus it requires empirical support. However, my main work in this chapter is not to rigorously defend the overestimation thesis with systematic empirical evidence. I have already provided characteristic examples of meta-research findings that reveal overestimation bias, and Jacob Stegenga (2018) provides countless other examples and considerations.

Meta-Evidentialism is mainly a methodological thesis. It is a prescription for evidence evaluation that instructs us to include meta-evidence in our evaluation of therapeutic evidence, which will generally change our beliefs about the evidence and the effect size. The qualification that we should "generally" lower our confidence builds into the thesis the possibility of defeaters. By default, we should modify our confidence—unless we have reason not to do so. The default is defeasible, subject to defeat, when we have reason to maintain our prior confidence in the effect estimate despite the meta-evidence. Defeaters are reasons to think that the trends detected in the meta-research are not a threat to the particular piece of therapeutic evidence on the table. Modifying our confidence after learning of the meta evidence should be the default (rather than *not* modifying our confidence) because meta-evidence is relevant for assessing our confidence in the effect estimate, as I will argue in the next section. It often goes beyond the considerations included in EBM "risk-of-bias" tools or evidence grading frameworks.

We can illustrate the meta-evidentialist's approach through the example of industry bias. The meta-research by Lundh et al. (2017) found that the ratio of the frequency of positive efficacy results in industry-funded studies to the frequency of positive efficacy results in non-industry-funded studies is 1.27. The most plausible interpretation of this finding is the standard one: that this result is the manifestation of an "industry bias" (Lundh et al. 2017). The meta-research warrants generally having a lowered confidence in the results of industry-sponsored therapeutic evidence compared to our confidence in non-industry-funded evidence (and lower than our confidence before

learning of the research on industry bias). However, the meta-evidence can be defeated if we have reason to believe that the results of a particular industry-sponsored study are less prone to bias (e.g., if special safeguards are in place to prevent manipulation or "spinning" of the study results or if there is little incentive to do so given the unique circumstances).[19]

The meta-evidentialist should then revise their initial confidence. But by what degree? Surely it is possible to under-correct or to over-correct, to underestimate or to overestimate the magnitude of the problem. Indeed, many people who are aware of the problem of industry bias might already be adjusting their confidence in industry-supported studies in a way that is poorly calibrated to the meta-research. The meta-evidentialist can look to the meta-research findings themselves for guidance. For instance, the results of Lundh et al. (2017) can be used to derive a simple correction factor or "bias coefficient" (Erasmus 2023) for our confidence in the effectiveness of the medical intervention. If we assume that the true frequency of effective medical interventions is the same in industry-funded compared to non-industry-funded research and that the ratio of 1.27 cited above results entirely from an increased frequency of false positive findings among industry-sponsored studies (due to bias), then the meta-evidentialist can multiply their unadjusted credence in the effectiveness of an intervention by the inverse of this ratio (1/1.27) to calculate a corrected credence that adjusts for the influence of industry funding on a study's results. (Solomon 2020a makes a similar suggestion.)

One might worry that using such a quantitative bias coefficient is an act of faux quantitative precision. After all, the industry sponsorship meta-analysis is not designed to provide a "confidence correction factor," and a clinician is typically not able to put a precise number to their unadjusted confidence anyway. However, its function is not necessarily to provide us with a precise adjusted credence, but rather to help us conceptualize roughly how much of a cognitive impact the meta-research evidence should have. (Our corrected confidence should be somewhere between "unchanged" and "half as great as

[19] Resnik and Elliott discuss five factors relevant to determining the likely influence of industry funding on the reliability of study results: "whether sponsors, institutions, or researchers have a significant financial stake in the outcome of a study; whether the financial interests of the sponsors, institutions, or researchers coincide with the goal of conducting research that is objective and reliable; whether the sponsor, institution, or researchers have a history of biasing research in order to promote their financial goals; how easy it is to manipulate the research in order to achieve financial goals; and whether oversight mechanisms are in place which are designed to minimize bias" (2013, 184). See also Elliott (2014).

before.") On the other hand, precise correction factors may be useful for systematic reviewers who are trying to formally weight the results of different studies in their analysis.[20] That said, I offer my industry bias coefficient mainly as a proof of concept. (See J. Stegenga 2018 and Erasmus 2023 for more general models for representing the influence of bias on study results.) Further work is needed to derive more sophisticated analytic approaches to correcting therapeutic estimates (or our confidence in those estimates) using meta-research findings.

Another promising concrete proposal is to include industry sponsorship (or potential conflicts of interest more generally) as an item in risk-of-bias tools or in evidence hierarchies.[21] For instance, in the GRADE approach to evidence appraisal, the evidence user first forms an antecedent confidence ("high" or "low") based on the design of the study (randomized or observational, respectively). They then upgrade or downgrade their confidence based on factors such as "indirectness" of the study comparison with respect to real-world comparisons or the threat of publication bias. Industry sponsorship could be added to this list, prompting users to assess for risk factors, mechanisms, or symptoms of industry bias and determine its likelihood in a given case.[22] While industry sponsorship generally incurs a *risk* of bias and does not guarantee that a given industry-funded study is biased, the same applies for other GRADE items—even lack of randomization in the study design. A more serious limitation of this proposal is that industry bias may be variously due to indirectness, publication bias, or other independently listed factors, so including industry sponsorship in GRADE's list may result

[20] One might also worry about heterogeneity in the results of studies included in the Lundh et al. (2017) meta-analysis: the magnitude of association between industry sponsorship and trial outcome varies considerably among included studies. Narrower analyses exist that measure the association for a specific medical specialty (e.g., Djulbegovic et al. 2013), disease (e.g., Djulbegovic et al. 2000), or treatment (e.g., Djulbegovic et al. 1999). Lundh et al. (2017) also measure the association separately for studies at higher risk of bias versus those at lower risk of bias according to a standard risk-of-bias tool. These kinds of results might be used to derive more specific and informative correction factors.

[21] This idea has also been proposed by Shea et al. (2007), Every-Palmer and Howick (2014), Lundh et al. (2017), and Solomon (2015, 2020a). One challenge with including industry bias in risk-of-bias tools is that those tools are generally applied to the study itself, while industry bias could result from lack of external validity, which depends partly on the target population. As suggested in section 8.2, there are several ways to get around this problem, including assessing external validity for typical target populations.

[22] In their review of 1,567 interventions included in the Cochrane Database of Systematic Reviews, Jeremy Howick and colleagues (2022) found that 58% of interventions had a highest GRADE rating of "low" or "very low," corresponding to a weak confidence in the effect estimate. My specific proposal here suggests that some interventions should receive an even lower rating. However, Howick et al.'s findings do not entail that most of the time we should have low (or very low) confidence in effectiveness (meaning: a net positive effect) because a low confidence in an effect estimate is consistent with believing that the true effect is positive, just smaller.

in "double counting" the influence of these forces and over-adjusting our confidence.

Proposals like these serve to illustrate the ways in which Meta-Evidentialism transcends traditional EBM by relying on meta-research evidence to directly inform evidence evaluation. The authors of the *Cochrane Handbook* suggest that "[w]here possible, assessments of risk of bias in a systematic review should be informed by evidence," particularly meta-epidemiologic research. However, in disagreement with my proposal to use meta-epidemiologic research to develop bias coefficients, they "do not encourage use of bias adjustments in the context of a Cochrane Review." Moreover, they "discourage the inclusion of conflicts of interest *directly* in the risk-of-bias assessment" (Boutron et al. 2024), while allowing that conflict-of-interest assessments can inform the application of risk-of-bias tools and subgroup analyses.

Medical Nihilism and Medical Meta-Evidentialism are not incompatible; one could be a nihilist *and* a meta-evidentialist. In fact, Medical Nihilism (we should have low confidence in the effectiveness of medical interventions), coupled with the assumption that our current confidence is higher than low, implies that we should lower our confidence, as Meta-Evidentialism recommends. The medical nihilist believes that our confidence in effectiveness needs so much downward adjustment that our corrected confidence winds up being low. However, one could be a meta-evidentialist but not a nihilist by holding that we should lower our confidence in effectiveness, but not to the extent that we come to have low confidence in the effectiveness of medical interventions. That said, Meta-Evidentialism is a superior response to the problems facing medical research because it is not vulnerable to the objections to Nihilism advanced in the previous section.

First, Meta-Evidentialism dodges the broad-brush objection because it does not assign an average confidence to our collection of medical interventions. It recognizes that our medical interventions are a mixed bag by starting with a piece of therapeutic evidence before correcting for threats to internal and external validity. On the other hand, one might worry that applying meta-research findings to first-order evidence is to use too broad a brush. For instance, Mercuri (2018) argues against a categorical downgrading of confidence in effectiveness due to widespread publication bias and other biases, and Broadbent (2019) argues that we should apply evidence specific to the particular case rather than evidence of general trends in medical research. However, the adjustment procedure that Meta-Evidentialism

recommends is defeasible and can be forgone if there are reasons to believe that the particular therapeutic evidence bucks the trend measured in the meta-research. At the same time, it is because biases like publication bias and industry bias are tricky to detect and correct in the first-order evidence[23] that we must apply evidence of the general trend by default.

Next, the multiple measures of effectiveness objection applies only when we are forced to draw a line between an effective and an ineffective intervention based on the magnitude of effect because there are multiple ways of quantifying effectiveness and they don't always agree on whether an effect is clinically important. Meta-Evidentialism concerns our confidence in the effect estimate itself, whatever measure of effectiveness is used. The advantage of focusing on the effect estimate is that it is central to informing uncertainty in medicine, as I argued in Chapter 6. Granted, this effect estimate can simply be "the effect is net positive" or "the intervention is effective." However, for Meta-Evidentialism, "effective" means that the effect is greater than zero rather than that the effect is greater than "small," which avoids the problem of choosing a way of quantifying effectiveness and then using it to draw a line between small and larger effects.

Finally, the vague-probabilities objection protests the use of vague terms like "low" and "high" to describe probabilities. Meta-Evidentialism claims that we should lower our confidence, which has a clearer meaning. (How much we should lower our confidence depends on the token first-order evidence and the token meta-evidence.) However, in the previous section I also suggested that meta-research evidence does not allow us to estimate the probabilities in the Master Argument for Medical Nihilism. Is meta-research of any use to Meta-Evidentialism? The answer is that Meta-Evidentialism makes better use of meta-research evidence because it starts with first-order therapeutic evidence and then makes adjustments from there. Our initial confidence anchors our adjusted confidence. Moreover, meta-research on biases in therapeutic research is often comparative: it measures how much bias exists among certain kinds of studies through comparison with other kinds of studies. This can sometimes allow us to make a correction to a given evidence token of one of these kinds (as I demonstrated through the example of industry bias). Moreover, even when meta-research evidence does not allow us to make quantitative corrections, all Meta-Evidentialism demands is

[23] The difficulty is noted and discussed in, for instance, Lau et al. (2006), Guyatt, Oxman, Montori et al. (2011), S. Goodman and Dickersin (2011), Resnik and Elliott (2013), and Lundh et al. (2017).

that we lower our confidence in an effect estimate by some amount, without requiring that we precisely quantify our corrected confidence.

Medical Meta-Evidentialism advises us to lower our confidence in effect estimates (and raise our confidence in lesser effect estimates) after learning of relevant meta-evidence. Rigorously developed tools are needed to facilitate these adjustments, and the field of meta-research can provide an evidence base for such tools. In the next section, I provide a positive argument for Medical Meta-Evidentialism.

## 8.5. The Argument from Meta-Research Evidence

The *argument from meta-research evidence* defends both the premise that meta-research findings have evidential bearing on medical hypotheses and the premise that we should respect this evidential bearing. When we do so, we find ourselves adjusting our confidence in medical hypotheses about the therapeutic effect, as Medical Meta-Evidentialism recommends. I develop this argument, while defending it against three objections: an irrelevance objection, a screening-off objection, and an underdetermination objection.

According to the argument theory of medical evidence, medical evidence supports premises in inferences or arguments for medical hypotheses, which includes hypotheses about a patient's response to treatment. The support that medical evidence provides is often indirect, supporting a premise in a sub-argument for a premise in the main argument. Even the results of clinical trials or meta-analyses of medical interventions—clear-cut examples of medical evidence—support a medical hypothesis only indirectly by showing that the intervention is efficacious in a study population, an early premise in an overall chain of reasoning terminating in the hypothesis that a patient will respond to treatment in a certain way. Meta-research findings also provide indirect evidence for medical hypotheses and equally should be viewed as medical evidence.

To understand the relevance of meta-research findings, we can distinguish first-order evidence from meta-evidence. *First-order evidence* is direct evidence for a hypothesis such as a positive result in a clinical trial, which is direct evidence for the hypothesis that the treatment is efficacious for the trial population because it supports a premise in a causal inference about the trial population. Typically when we speak of medical evidence in EBM, we have some first-order evidence in mind. In comparison, *meta-evidence* is evidence relevant to

evaluating the evidential support that the first-order evidence lends to the hypothesis. Meta-research findings about bias in the first-order evidence base are meta-evidence because bias is relevant to the support that first-order medical evidence lends to efficacy or effectiveness hypotheses, and thus meta-research findings about bias are relevant to assessing first-order evidential support. We can more specifically refer to meta-evidence from meta-research as *meta-research evidence*, allowing that other kinds of meta-evidence might derive from different sources, such as an expert's personal experience with the quality of a certain medical literature.

A key premise in my argument from meta-research evidence is that meta-evidence is also (indirect) medical evidence: not only is it relevant to evaluating the first-order evidence's support for a hypothesis, but it is also relevant to the hypothesis itself. Say we have some first-order evidence in the form of a positive clinical trial or meta-analysis, in which the intervention appears efficacious compared to placebo. Upon learning about meta-research evidence on publication bias (for instance), our confidence in the efficacy hypothesis should be diminished because the meta-research evidence lends support to the alternative hypothesis that bias explains the first-order study results. Because bias is a competing hypothesis, it robs support from the efficacy hypothesis, which subsequently lowers support for the medical hypothesis that a patient will respond to the intervention in a way suggested by the trial results. On the argument theory, meta-evidence about bias is evidence *against* the efficacy hypothesis. More specifically, it at least partially undermines the support that first-order evidence provides (without rendering the first-order evidence irrelevant).

In recent research in epistemology, evidence about evidence, or evidence about an agent's reasoning, is often called "higher-order evidence," and it presents several interesting puzzles.[24] One such puzzle is that believing both first-order evidence as well as higher-order evidence that casts doubt on the

[24] See, for instance, Feldman (2005), Roush (2009), Christensen (2010), Horowitz (2014), Sliwa and Horowitz (2015). Because the paradigm examples of higher-order evidence involve either the reliability of an individual epistemic agent or the correctness of their inferences from the first-order evidence, I will not assume that meta-research findings are higher-order evidence and will continue to use the uncommitted term "meta-evidence" in describing them. Meta-research evidence is not evidence about a particular agent's reasoning but rather about public evidence from which many agents reason. Moreover, meta-research is not straightforwardly about the accuracy of inferences from first-order evidence. A paradigm example of higher-order evidence is the existence of expert disagreement about a hypothesis (Feldman 2005). Daniel Friedman and Dunja Šešelja (2023) discuss the role of higher-order evidence during fast science, using the example of scientific expert disagreement during the early months of the COVID-19 pandemic over the importance of aerosol transmission of SARS-CoV-2.

reliability of the first-order evidence amounts to believing something like "H, but my evidence doesn't support H" (Feldman 2005; Horowitz 2014). There is a similar tension between meta-evidence and first-order therapeutic evidence. The first-order evidence supports the first-order hypothesis with a certain level of confidence, while the meta-evidence is evidence that the first-order evidence should support a different level of confidence. One way out of this tension is to drop our confidence in the first-order hypothesis in line with the higher-order information, which is the solution advocated by Meta-Evidentialism.

However, an *irrelevance objection* might argue that in lowering our confidence in the first-order hypothesis, we fail to respect the rational bearing of the first-order evidence. Let's assume we correctly judge our first-order evidence and its relevance for the hypothesis. The irrelevance objection holds that for us to change our mind after learning about the meta-research is irrational because our new beliefs are no longer a good reflection of our first-order therapeutic evidence, which, after all, is indisputably relevant to the hypothesis. If either the first-order evidence or meta-evidence must trump the other, the first-order evidence should surely win. One way to develop this objection further is to consider through what inference therapeutic study results support the hypothesis. The hypothesis in question concerns the efficacy or effect of a treatment, which are causal concepts. We must therefore decide the hypothesis through a causal inference. Notably, our best theories of causal inference do not seem to rely on meta-research, but they *do* accommodate clinical trials and other first-order studies.

I think that this objection could call on any popular theory of causal inference, including those that use causal Bayes nets, counterfactual frameworks, or the probabilistic theory of causality. We can represent causal inference from trial results or other epidemiologic study results using the following highly abstracted argument scheme:

1. Metaphysical assumption(s).
2. Difference in outcome or probability of outcome between groups (from first-order evidence).
3. The right kind of causal comparability between groups (in first-order evidence).

H: The intervention caused a difference in outcome or probability of outcome.

Different theories formalize the premises differently. For instance, Nancy Cartwright's (2010) theory of the ideal RCT uses the probabilistic theory of causality for the metaphysical assumption and understands causal comparability to mean an equal distribution of causally homogeneous subpopulations. My own theory of the ideal comparative group study (Fuller 2019) formulates causal comparability as an equal contribution of a variable "C," representing all the complex causal conditions.

However we choose to fill in the details, we need only first-order study results to support the second premise, and we can presumably satisfy the third premise by checking that there are no systematic differences between study groups. Thus the first-order study results are evidence for the second premise, and evidence for the third premise can be generated through traditional critical appraisal by checking the first-order evidence for evidence of bias. Regardless of what the meta-research shows, the first-order evidence continues to demonstrate a difference in outcome, so the meta-research does not affect premise 2—it does not lead us to bracket the support provided by the first-order evidence (Christensen 2010). Further, the premises are engineered to deductively imply the conclusion H (Cartwright 2010); thus meta-research seemingly has no effect on the argument's validity—it is not an "undercutting defeater" for the argument from first-order evidence (Feldman 2005). With all that in mind, is the meta-research relevant at all to this causal inference?

My answer is that the function of meta-research evidence is to rationally revise our confidence in the first-order evidence and our confidence in the premises of our argument, but not necessarily in the same premise that the first-order evidence supports. In this example, meta-research on publication bias should lower our confidence that the first-order evidence is unbiased by lowering our confidence that the study groups are causally comparable in the right ways, *premise 3*.[25] The meta-research evidence is evidence against premise 3 and thus against hypothesis H. It may not provide enough evidence against premise 3 that we should accept that premise 3 is false and that our causal inference is unsound. Nonetheless, it weakens the strength of the first-order evidence or the support it provides to the hypothesis via this

[25] Publication bias is relevant to assessing causal comparability in a meta-analysis because selective publication of positive study results can misleadingly trim the statistical variability in the sample so that statistical analysis is biased. Because meta-research evidence on industry sponsorship may reveal bias in the wider sense, relevant to internal validity and/or external validity, it should lower our confidence in the conjunction of premise 3 of this causal inference and the sufficient similarity assumption required by a Generalization inference (see Chapters 5 and 7).

argument because premise 2 (reporting the first-order evidence) warrants the hypothesis only in conjunction with premise 3. A difference in outcome in a trial warrants the conclusion that the intervention caused the difference in outcome only if the study groups are comparable. By altering our confidence in first-order evidence and in premise 3, meta-evidence alters our confidence in the hypothesis that the first-order evidence supports,[26] in line with Meta-Evidentialism. Because the arguments for which meta-evidence is relevant are ultimately arguments (or sub-arguments) for medical hypotheses, meta-research evidence is medical evidence according to the argument theory of medical evidence.

Meta-research evidence is evidence of bias if we adopt the inferential perspective on bias introduced in section 8.2, in which biased evidence is misleading evidence, evidence which supports an unsound medical inference. While the first-order therapeutic evidence supports a causal inference by establishing premise 2, meta-research evidence suggests that the causal inference is unsound by providing evidence against premise 3. The first-order evidence may be veridical and truly support this inference, but if premise 3 is false then the causal inference is unsound and the first-order evidence's support for the inference and its conclusion is misleading. Meta-research evidence provides evidence that this is in fact the case.

The first component of my argument from meta-research evidence is that meta-research evidence is medical evidence, which I've now established. The second component of my argument is that we should respect the medical evidence provided by meta-research and use it to adjust our confidence in therapeutic effect estimates provided by first-order research. In making my case, I could rely on a "principle of total evidence" (Carnap 1945): if certain meta-research findings are medical evidence for a therapeutic hypothesis, then on that basis alone they should influence our evaluation of the hypothesis, lest we fail to respect our total body of evidence. However, I do not need to rely on any grand philosophical principle here. Because meta-research evidence reveals bias in first-order evidence, meta-research evidence should influence

[26] We can formalize my argument using probabilities to represent our confidence. We can let $p(H)$ be our confidence in H, $p(H|E)$ be our confidence in H given positive (unbiased) first-order evidence, $p(H|\neg E)$ be our confidence in H absent the first-order evidence, $p(E)$ be our confidence that the (unbiased) evidence obtains, and $p(\neg E)$ be our confidence that the evidence does not obtain. On the total probability equation, $p(H) = p(H|E)p(E) + p(H|\neg E)p(\neg E)$. Because E confirms H, $p(H|E) > p(H|\neg E)$. Thus whenever meta-research evidence E' changes our confidence in E (the $p(E)$), we must update our confidence in H (the $p(H)$) to avoid being irrational.

our evaluation of the hypothesis because bias should influence our evaluation of the hypothesis. To suggest otherwise is to flaunt basic principles of evidence-based medicine, clinical epidemiology, and science. While EBM uses methodological criteria to assess for bias in first-order evidence, my proposal uses meta-research evidence. If our evaluation of medical hypotheses is to be "evidence-based," it must respect our meta-research evidence.

However, a further objection to the argument from meta-research evidence, a *screening-off objection*, worries that meta-research provides information about a broad and heterogeneous reference class of first-order evidence. If we had more specific information about whether the particular first-order evidence token was biased, then this assessment would screen off the meta-research evidence. More specific information can be provided by traditional EBM critical appraisal, by inspecting the specific evidence token (or so the objection goes).[27]

I anticipated this objection in the previous section when I noted that Meta-Evidentialism permits that we should not abide by the meta-research findings when we can better judge the workings (or absence) of bias in the particular piece of first-order evidence. We should lower our confidence in the therapeutic effect estimate by default but should also allow considerations suggesting that the particular first-order evidence is likely to defy the general meta-research trend to defeat this default. Unfortunately, various biases that meta-research investigates are difficult to detect in the individual case, which is why we need to rely on meta-research evidence by default as a proxy for the detection of bias in the particular first-order evidence token. In their systematic review of industry sponsorship, Lundh and colleagues (2017) conclude, "Our analyses suggest the existence of an industry bias that cannot be explained by standard 'Risk of bias' assessments" (2) and argue from their data that "an assessment of [industry] sponsorship should therefore be used as a proxy for these mechanisms" of industry bias (21). For publication bias, trial registries are incomplete (Hopewell et al. 2009) and often fail to alert us to unpublished studies. While graphical and statistical tests of publication bias may be useful, EBM's GRADE group concedes that they are error-prone and rely on strong assumptions, and that publication bias is "difficult to predict for individual systematic reviews" (Guyatt, Oxman, Montori, et al. 2011,

[27] See Tresker (2022) for a related objection regarding the epistemic relevance of a study author's financial conflict of interest disclosure.

1278). Therefore, meta-research evidence is not always screened off, and we should use it to adjust our confidence whenever it is not.

A final objection, an *underdetermination objection*, notes that there are many possible explanations for these meta-research findings beyond bias. For example, I have (following a consensus in the medical literature) assumed that the association between industry sponsorship and positive study results quantified by Lundh et al. (2017) is due to an industry bias (in the wider sense of "bias" that encompasses both internal and external validity). Two other possible interpretations are (1) that drug and device companies more often study effective treatments compared to non-industry-sponsored researchers (there is no bias), and (2) that non-industry-sponsored studies more often suffer from bias that leads them to systematically underestimate the effectiveness of drugs and devices ("non-industry bias" explains the association). Moreover, quantitative correction factors are difficult to develop from meta-research findings; even the simple one I developed for industry bias in the previous section requires strong assumptions. In short, meta-research findings underdetermine our confidence in the effectiveness of medical interventions.

My response to this objection is threefold. First, Meta-Evidentialism need not consider all meta-research findings to be meta-research evidence, only those we can judge are relevant to assessing the evidential import of first-order therapeutic evidence. Second, while there are usually several possible interpretations for meta-research findings on industry sponsorship, selective publication, and similar findings, when we allocate some credence to all the plausible interpretations, on net they should lead us to have lowered confidence in a positive effect estimate. For instance, interpretation (2) for the results of Lundh et al. (2017) is far less plausible than the existence of industry bias because while pharmaceutical industry tactics that produce hidden bias are widely documented,[28] mechanisms that might explain hidden non-industry bias are not widely documented. Even if we permit the plausibility of interpretation (1) and split our credence between this interpretation and the industry bias interpretation,[29] on net we should lower our

[28] See Rochon et al. (1994), Safer (2002), Lexchin et al. (2003), Sismondo (2008), Goldacre (2012), McGarity and Wagner (2012), Boutron et al. (2024).

[29] My own credence in the industry bias interpretation is higher. The industry bias interpretation receives additional support from further evidence summarized in Lundh et al. (2017). In a separate meta-analysis conducted by the authors, trials comparing drug A to drug B that were sponsored by the maker of drug A were 5.92 times as likely to favor drug A compared to trials of A versus B that were sponsored by the maker of drug B. In this analysis, the same set of drugs are studied in the company

confidence in a positive industry-funded study's results. Finally, even when meta-research does not suggest a quantitative correction factor, it often still suggests the direction in which adjustments should be made, allowing us to be more confident or less confident in the effect estimated by first-order evidence. This can be helpful when the decision about whether to use the intervention is balancing on a knife's edge and can be tipped in one direction or the other.[30]

In summary, not only are certain meta-research findings genuine medical evidence, but this evidence should be used to adjust our confidence in effect estimates provided by first-order evidence. In short, we ought to be meta-evidentialists.

## 8.6. Conclusion

Among various skeptical attitudes we could adopt toward contemporary therapeutic research, I argue for Medical Meta-Evidentialism over Medical Nihilism. Meta-research findings about the existence of biased or misleading evidence in the medical literature lend fuel to contemporary strains of therapeutic skepticism. According to an inferential perspective, medical evidence is misleading or biased when it supports a premise in an unsound medical inference. One recently defended therapeutic skepticism is Medical Nihilism, the thesis that we should have low confidence in the effectiveness of contemporary medical interventions. I rejected Medical Nihilism for three reasons: it is too broad, it depends on measures of effectiveness that do not fully capture the value of interventions, and it is supported by a vague probabilistic master argument. As an alternative response to the problem of biased and misleading evidence, I advanced Medical Meta-Evidentialism: we should generally lower our confidence in effect estimates provided by first-order therapeutic evidence (and raise our

A–sponsored trials versus the company B–sponsored trials, which rules out the no bias interpretation (which, again, posits that the industry sponsorship association is explained by different drugs being studied in the different arms of the meta-analysis). The "non-industry bias" interpretation is also not possible here because none of the trials is non-industry sponsored. Industry sponsorship–related bias seems the only plausible explanation. If industry bias can plausibly explain a 5.92-fold difference in this analysis, it can plausibly explain a 1.27-fold difference in the industry versus non-industry meta-analysis.

[30] See Fuller (2018a) for further elaboration of these objections and my responses.

confidence in lesser effect estimates) after learning of meta-evidence about the first-order evidence. Meta-Evidentialism is immune to my objections to Medical Nihilism. It is also supported by the argument from meta-research evidence: by providing evidence of bias in first-order therapeutic evidence, certain meta-research findings are medical evidence, and we should respect this medical evidence. As in previous chapters in Part 3 of this book, an argument theory of medical evidence proves its worth in the new modern medicine by clarifying important research concepts (such as bias), elucidating the relevance of diverse evidence, and, potentially, improving the evaluation and use of medical evidence.

# PART 4
# EPIDEMIOLOGICAL MEDICINE

# 9

# Twenty-First-Century Medicine

## Epidemiological, Personalized, or Postmodern?


### Summary

During the second half of the twentieth century, modern scientific medicine went epidemiological. An epidemiological medicine based in epidemiological science became the dominant model of medicine in society. How will twenty-first-century medicine compare? I argue that epidemiology continues to permeate all aspects of contemporary scientific medicine, and that far from being in decline, recent events and trends reveal new directions in which epidemiological medicine is developing, as it continues to function as the dominant model of scientific medicine.

In one revolutionary possible future, a personalized or precision medicine that gives scientific priority to the results of laboratory research in molecular biology becomes the new model of scientific medicine. Another possible future—revolutionary in a different way—is a postmodern medicine in which scientific medicine loses its societal dominance. A close look at the personalized/precision medicine movements instead reveals that their more realistic vision is of a new epidemiological medicine grounded in molecular epidemiology. Meanwhile, by analyzing two case studies—one of the pseudoscientific medicine homeopathy and its epidemiological research, the other of resistance to scientific medicine during the COVID-19 pandemic—we discover signs of epidemiological medicine's health overwhelming signs of its imminent deterioration. Medicine remains epidemiological. And yet several problems of the new modern medicine discussed in this book are shown here to result from the way that contemporary clinical medicine is modeled after epidemiology. Twenty-first-century medicine will be defined by how it responds to these problems.



*The New Modern Medicine*. Jonathan Fuller, Oxford University Press. © Oxford University Press 2025.
DOI: 10.1093/9780190066178.003.0010

## 9.1. The Epidemiological Revolution

In 1917, when he was twenty years old, Austin Bradford Hill developed pulmonary tuberculosis (Doll 1994). Despite germ theory, despite the theory of "internal antisepsis" from bacteriology and the idea that "magic bullets" for infections were waiting to be discovered, there was still no successful drug for that scourge of the nineteenth century. Hill was expected to die, and nearly did. After a nine-month course, however, he improbably survived. Weakened by the experience and convalescing at home, he forfeited his ambition of becoming a physician, which would have required attending university in person. For the same reason, he did not pursue a degree in natural science like his father, a physiologist, which would have required laboratory work. Had he done either, Hill might have been lost to laboratory medicine, the then-dominant tradition of modern scientific medicine.

Instead, Hill pursued economics and then statistics, becoming an epidemiologist (as history would remember him) at the new London School of Hygiene and Tropical Medicine, and later, director of the Statistical Research Unit for the British Medical Research Council (MRC). There he convinced the investigators of a clinical trial to implement randomization of participants to trial groups, an uncommon practice at the time. The resulting study is often remembered as the first modern randomized clinical trial. The treatment was streptomycin, one of the first antibiotic magic bullets; the disease: tuberculosis (MRC 1948). Nearly thirty years after TB changed the trajectory of his own career, Hill's randomized controlled trial (RCT) of streptomycin for tuberculosis set a new course for therapeutic evaluation in medicine. The RCT was later central to the new science of clinical epidemiology, which became integrated into medical practice by way of the evidence-based medicine (EBM) movement.

The same year that the results of the streptomycin trial were published, Hill pulled off a second trick. While mortality from tuberculosis and other infections was waning, death from cancer and coronary artery disease, noncommunicable diseases, was waxing. Asked by the MRC to investigate the cause of the rising rate of lung cancer, Hill began a case-control study of inpatients with lung cancer that implicated smoking (Doll and Hill 1950). He then undertook a prospective cohort study of British doctors that established an association between smoking and lung cancer (Doll and Hill 1954), and eventually between smoking and other ailments such as coronary

artery disease. Based on principles he articulated for inferring causation from correlation in epidemiologic studies, Hill concluded that smoking caused lung cancer—a contribution to preventive rather than therapeutic medicine. Hill's (1965) principles became general guidelines for inferring causation in epidemiology, an analogue of Koch's postulates in bacteriology. The study designs and principles that Hill wielded against cancer helped establish noncommunicable disease epidemiology as a public health science as well as a basic science for medicine.

And so epidemiology became established in medical research and health governance as well as integrated into medical education and practice. This was the new modern medicine: epidemiological medicine. It is shaped by new disease epidemics like cancer and heart disease but also new norms of scientific evidence and the new sciences of epidemiology. In Part 1 of this book, I argued that modern medicine began with scientific medicine, in which clinical medicine is integrated with and modeled after modern medical sciences, and in which scientific medicine occupies a unique position of influence in society. While laboratory science had greater clinical priority in the old medical model of laboratory medicine, from the middle of the twentieth century epidemiological science has come to enjoy a privileged place in the new medical model of epidemiological medicine, as medicine underwent a multilayered epidemiologic transition. This transition involved not only changing population patterns of disease and the rise in mortality from noncommunicable diseases (the classic "epidemiologic transition") but also changing ways of understanding disease as well as evidence. In Part 2, I investigated the theories, causes, and nature of medicine's new epidemic diseases, including cancers and chronic diseases. In Part 3, I examined medical treatment, risk, and inference in a medicine "based" in epidemiological evidence.

In recent decades, genomics and other -omics sciences have enjoyed increasing medical research support, including publicly funded research at agencies like the MRC and the National Institutes of Health (NIH; Tabery 2023). This has led some to prophesy a coming revolution in medicine—"personalized" or "precision medicine"—in which medicine is more closely aligned with these newer sciences. Some of the more radical models of personalized medicine involve making inferences from molecular biological theory, computer modeling, or in vitro experiments. Laboratory experiments and laboratory science reasoning were central to the laboratory medicine model of scientific medicine that preceded epidemiological

medicine. Does personalized or precision medicine represent a new laboratory medicine and a *newer* modern medicine?

On the other hand, some have worried about the recent rise of a powerful anti-science movement threatening medicine (Hotez 2023), in which scientific medicine is weakened while pseudoscientific medicine flourishes. The recent COVID-19 pandemic is a promising case study of public resistance to scientific medicine and its potential downfall. Meanwhile, the emergence of meta-analyses of homeopathy RCTs seeming to support homeopathic medicine is a promising case study of the rise of pseudoscientific medicine. If indeed there are growing cracks in scientific medicine's once-solid edifice, it could portend a period in which scientific medicine no longer enjoys unique societal dominance, a *postmodern* medicine.

In this final chapter, comprising Part 4 of the book, I consider whether the picture of the new modern medicine developed in the rest of the book is becoming outdated. However, I conclude that epidemiology continues to permeate all aspects of contemporary scientific medicine, and that far from being in decline, precision medicine and the COVID-19 pandemic reveal new directions in which epidemiological medicine is developing, as it continues to function as the dominant model of scientific medicine. In section 9.2, I review epidemiological medicine and the ways that epidemiology has shaped the problems explored in this book, including seven problems of the new modern medicine. In section 9.3, I clarify two visions of personalized or precision medicine, arguing that the more realistic vision is of a molecular epidemiological medicine instead of a new laboratory medicine. Then in section 9.4, I argue that rather than heralding medicine's postmodern decline, homeopathy RCTs provide a surmountable challenge to epidemiological medicine's scientific credentials, while the COVID-19 pandemic showcased the power that epidemiology wields in mainstream medicine today. Though in both cases, faults in the foundation of the new modern medicine are exposed.

## 9.2. Epidemiological Medicine

In the twentieth century, epidemiological medicine emerged as the new model of scientific medicine, reshaping medicine's perspective on disease and evidence. In the process, it created several problems for the new modern medicine. Throughout the book, I examined some of these problems. Here I summarize how they result from the role of epidemiology in medicine.

By the early 1900s, clinical medicine was dominated by a model of scientific medicine known as laboratory medicine, in which clinical medicine was modeled after experimental laboratory sciences. In the mid-1900s, medicine's threefold epidemiologic transition in disease, evidence, and science transformed it into the epidemiological medicine of today. To review, in the epidemiological medicine model, the etiology of diseases is understood to be molecular and multifactorial. Genes and other risk factors help explain noncommunicable diseases of the epidemiologic transition, especially cancer and heart disease. Diseases are often understood to be chronic and asymptomatic, with nonspecific pathology. Doctors prevent, treat, and manage these diseases, particularly by reducing the risk of their occurrence and of their manifestations. Randomized controlled trials are the best evidence for guiding medical management, and predictions are made about treatment success by extrapolating from their results—that is, after critical assessments are made of their validity and whether they are biased.

In earlier chapters, I discussed ways that various epidemiological sciences have influenced these features of epidemiological medicine, especially noncommunicable disease epidemiology and clinical epidemiology, but also molecular epidemiology and metascience. In the introduction, I also listed seven problems of the new modern medicine: molecular reductionism; multifactorial etiology; chronicity, asymptomaticity, and variable pathology; evidence hierarchies; meaning of medical risk; external validity of clinical trials; and biased evidence. Throughout the book, I examined these problems and often proposed solutions. Here I review how they are at least partly created by the influence of epidemiology on clinical medicine in order to show how epidemiology has so drastically molded the new modern medicine in its image in ways that are more and less obvious.

*Molecular reductionism* occurs when sub-personal explanations for a disease are given that are populated with molecular causes, including genes (Chapter 2). Reductionist explanations may fail to capture the complexity of diseases like cancer in which many causes interact in large multiscale networks, while also ignoring other important causes, including certain environmental and social causes. Although molecular biology has contributed greatly to our knowledge of the molecular etiology of diseases, epidemiology has also played a major part by providing data and mathematical modeling that supported a role for genetic mutations in carcinogenesis. The data included research in vital statistics showing that cancer increases with age across different types of adult cancer. The modeling included a multistage

theory that in the 1950s unified these observations by showing that they can be explained if cancer arises through multiple rate-limited steps, which were hypothesized to be genetic mutations. Today "genetic" or "molecular" epidemiology uses epidemiologic methods to study the genetic and molecular causes of disease in humans (Seyerle and Avery 2013), including specific genes involved in cancer. The limitations of reductive molecular explanations are thus to some extent limitations of epidemiological research into the genetic and molecular causes of disease. The theory that genetic changes cause cancer, born of both the laboratory and epidemiology, can focus research attention on genetic and molecular causes, diverting resources away from other research programs on cancer, that fatal disease of the epidemiologic transition.

All diseases, including infectious diseases, have *multifactorial etiology* (Chapter 3). The idea that infectious diseases are *mono*causal was imparted to medicine by its laboratory methods for isolating infectious germs and recapitulating human disease using animal models, as well as the practice of defining particular kinds of infectious disease according to a particular kind of germ. The contrasting idea—now pervasive in medicine—that all diseases (particularly noncommunicable diseases but also infectious diseases) have multifactorial etiology is not an obvious fact about their causation but a theory given to medicine by "risk factor epidemiology." It was a way for noncommunicable disease epidemiologists to make sense of the multiple associations between diseases and "risk factors"—factors that are statistically predictive of disease occurrence—revealed by their epidemiologic methods (case-control studies and cohort studies), particularly in studying the causes of heart disease and cancer. The multifactorial theory was soon crystallized in the form of the epidemiologist's sufficient-component cause theory, which holds that etiologic factors for a disease are component causes of variable sufficient causes of that disease. Risk factors are thought to causally interact in the multifactorial theory, just as they are made to statistically interact in epidemiologic studies through multivariate modeling. Multifactorial etiology provides a more complicated perspective on disease causation than monocausal etiology, while promoting multipronged approaches to prevention that target multiple risk factors, none of which is necessary or sufficient for the occurrence of the disease. That said, multifactorial etiology does not exclude the possibility of discovering universal etiology, even for noncommunicable diseases, which can sometimes be achieved by identifying specific kinds of disease with their specific pathology.

Many conditions of the epidemiologic transition, particularly noncommunicable diseases and mental disorders, are characterized by *chronicity, asymptomaticity, and variable pathology* (Chapter 4). Chronic conditions have a long duration and often a slow progression. They are chronic partly because they are incurable, due to a lack of laboratory breakthroughs resulting in cures. While acute infectious diseases can be studied through animal experiments lasting days or weeks, the long duration, slow progression, and delayed manifestations of chronic noncommunicable conditions are often best observed in epidemiologic studies with a long-term outlook, including backward-looking (retrospective) and forward-looking (prospective) studies. Next, while the asymptomaticity of many contemporary diseases has several origins, one of them is the recent pathologizing tendency to classify modifiable risk factors such as dyslipidemia, borderline hypertension, and non-insulin-dependent diabetes as diseases once medical interventions are found to reduce the risk of the overtly harmful diseases that they cause. Because risk and risk reduction are measured in epidemiological studies, epidemiology makes risk factor diseases possible. Finally, the *nonspecific pathology* of many contemporary disorders results from the relinquishing of the requirement of specific pathology that attended physiology- and bacteriology-based laboratory medicine. Unfortunately, nonspecific pathology can result in substantial heterogeneity in clinical presentation and treatment effect, which demonstrates the ongoing importance of investigating specific pathology and the laboratory sciences that do the investigating.

*Evidence hierarchies* are evidence grading methods produced by the EBM movement (Chapter 5). They inherit EBM's favoritism toward epidemiological evidence (especially RCTs) over the results of laboratory studies and pathophysiologic/mechanistic reasoning for medical decision-making. Sometimes hierarchies exclude non-epidemiological evidence altogether. They grade evidence using criteria that mainly apply to epidemiologic studies and identify threats to their accuracy such as bias and lack of external validity. Moreover, the main method of synthesizing evidence in EBM is meta-analysis, which by pooling effect estimates can combine the results of similar epidemiologic studies, but not with diverse evidence beyond epidemiology. Meta-analyses of RCTs are often found near the top of the hierarchy. Meta-analyses and RCTs are key methods of clinical epidemiology; indeed, EBM and its evidence hierarchies were designed by clinical epidemiologists. The EBM hierarchies view medical

evidence as measuring clinical outcomes, mirroring an emphasis within clinical epidemiology (and epidemiology more broadly) on quantification and measurement. Under EBM, "evidence" became virtually synonymous with the results of clinical epidemiology studies, the "evidence base" for clinical medicine. While they are useful, the evidence hierarchies suffer problems, such as a tendency to underappreciate the diverse roles that evidence plays in medical reasoning. To correct their clinical epidemiology/measurement-centric view of evidence, medicine could benefit from a theory in which any research data can serve as evidence—including research beyond epidemiological studies—when it supports a premise in a web of inferences or arguments.

The *meaning of medical risk* is a problem created by the attempt to make sense of epidemiological risk and risk reduction for clinical medicine (Chapter 6). Epidemiologic studies quantify medical prognosis using the population absolute risk and treatment effects using the population risk reduction, which measure the relative frequency of outcomes in a population and the impact of an intervention on that frequency, respectively. These epidemiologic measures are then the basis for ascribing individual risk and risk reduction to patients and for a new preventive medicine that deploys preventive interventions that reduce the risk of harmful conditions in populations and patients. Patient risk and risk reduction reflect an individual's tendency to have the outcome and the effect of an intervention in modifying that tendency. The concept of the patient's individual risk and the goal of reducing it reflect the attempt to translate epidemiological concepts and endpoints into the language and ends of clinical medicine. The problem of the meaning of medical risk is the challenge of determining what relevance group outcomes can have for an individual patient, what it means to ascribe risk and risk reduction to individuals, and whether epidemiological research measures individual patient risks at all (and if it doesn't, what the consequences are of inferring patient risks from population risks). Given that the population risk and risk reduction measured in epidemiological studies aggregate over and wash away any heterogeneity in outcomes among members of the population, they may not accurately predict the patient's individual risk, which can thus mislead doctors and patients. While epidemiological risks do have relevance for clinical medicine, the emphasis on patient-level risks should be replaced by a focus on medical uncertainty that can be informed by the results of epidemiological research.

Finally, the problems of the *external validity of clinical trials* (Chapter 7) and *biased evidence* (Chapter 8) have similar origins: they are problems with the use of epidemiological study results as evidence for medical decision-making. While similar problems can plague non-epidemiological evidence, the precise way they manifest in epidemiology and consequently in epidemiological medicine is specific to the design of its studies and the context of its research. The external validity of a clinical trial refers to the extent to which trial results predict the result of intervening in a target population. Many clinical trials may have poor external validity. The problem of external validity arises from the standard practice of making predictions about treatment effects by directly extrapolating from the results of clinical trials, rather than by relying on mechanistic reasoning or predicting from pathophysiologic theory, for instance. The challenge of assessing the external validity of a clinical trial can be partially overcome by making use of information about relevant causal mechanisms and biological structure, which is often provided by research outside of epidemiology.

Meanwhile, bias in contemporary medical evidence refers to systematic deviations from the truth in epidemiological study results that may then mislead doctors and patients. Bias can also describe the unsound inferences made from this misleading medical evidence about populations or individual patients. Metascience, which studies trends in epidemiological research, has documented many sources of bias, including biases that are difficult to detect when traditionally appraising epidemiological evidence, such as bias arising from industry sponsorship of research. Metascience itself often relies on epidemiology's methods for reviewing and analyzing the medical literature. To better handle the problem of biased evidence, evidence evaluation schemes can use the results of meta-research studies to account for hidden bias—yet another way that diverse evidence beyond traditional epidemiology has an important role to play in contemporary scientific medicine.

In summary, many of the problems facing scientific medicine in recent decades result at least partly from the challenges of integrating clinical medicine with the results of epidemiological research and modeling clinical medicine after epidemiology. In other words, they are special problems for an epidemiological medicine. In the next two sections, I consider whether epidemiological medicine is still an accurate vision of medicine's present and future, or whether current trends in science and society are remaking modern medicine yet again.

## 9.3. Personalized Medicine, Precision Medicine

In 2015, on the heels of a just-announced major federal research program, Francis Collins, director of the NIH, and Harold Varmus, director of the National Cancer Institute, prophesied "a new era of precision medicine" (2015, 793). The research program was the Precision Medicine Initiative, which seeks to harness progress in molecular research to make headway on understanding and treating cancers and, eventually, all diseases. The benefactor of this research, "precision" or "personalized" medicine, is often described as a medicine tailored to the individual: "the right treatment for the right patient." The language of a revolution in medicine swirls around precision medicine (National Research Council [NRC] 2011a), prompting one to wonder whether precision medicine represents a new model of scientific medicine, a *newer* modern medicine. Perhaps epidemiological medicine of the late twentieth century will soon be—or has already been—replaced. While there are truly revolutionary proposals for a radical personalized medicine, the leading vision in current research is for a more conservative precision medicine that represents a development within epidemiological medicine toward molecular epidemiology-based practice.

Epidemiological medicine is a species of scientific medicine, or science-based clinical medicine, and epidemiology is the scientific research in which it is characteristically grounded. Similarly, we can distinguish precision medicine as a variety of scientific medicine from precision medicine *research*. Much of the precision medicine literature and precision medicine initiatives are about precision medicine research. Precision medicine is usually described (especially by its proponents) as a vision for a future clinical medicine waiting just over the horizon, one we can reach through continued investment in precision medicine research (Hamburg and Collins 2010; F. Collins and Varmus 2015). Precision medicine research has not yet borne its promised fruits for most of clinical medicine, and thus clinical medicine is neither fully integrated with nor fully modeled after precision medicine research at present. In other words, precision medicine, as a model of scientific medicine, is currently mainly an aspiration rather than a reality (Plutynski 2022). Nonetheless, a few themes and exemplars recur in the literature on precision medicine, and I will use them to compare precision medicine as a vision for scientific medicine's future with epidemiological medicine, scientific medicine's present.

One can identify two visions in the precision/personalized medicine literature, one more conservative and the other more disruptive. In the more conservative vision, new medical categories that group patients into narrower bundles are used to make diagnostic, prognostic, and therapeutic inferences about particular patients. In the more disruptive or revolutionary vision, inferences are not made based on medical categories or from other patients at all but are unique to each individual patient. The terms "precision medicine" and "personalized medicine" are often used interchangeably, despite assorted attempts to assign a personalized and precise meaning to each (Tabery 2023). "Personalized medicine" is the older term and the one more commonly used in Europe, while "precision medicine" is the newer term popularized in North America by the 2011 NRC report *Toward Precision Medicine*. The authors of the report preferred "precision medicine" over "personalized medicine" because they worried that "personalized medicine" implied the more disruptive vision described above, while the report promotes the more conservative vision. Despite the rough equivalence between the two terms, it does seem that the term "personalized medicine" is more often applied to the more disruptive vision, while "precision medicine" is generally reserved for the conservative vision (Huang and Hood 2019). Thus, for ease of comparison, I will refer to the conservative vision as "precision medicine" and the disruptive vision as "radical personalized medicine."

A further distinction can be made between broad and narrow forms of either precision medicine or radical personalized medicine. In the broad form, any characteristic can be used to tailor prediction and intervention, while in the narrow form, biomarkers, especially genomic or other -omic molecular markers (transcriptomic, metabolomic, etc.), are used specifically (Gamma 2017; Letai 2017). While the broad form is usually touted as the eventual aim of precision or personalized medicine, most precision medicine research and examples of precision medicine are of the narrow form, hoping to take advantage of new scientific knowledge, resources, and methods developed in recent decades (Gamma 2013, 2017; Tabery 2023, 2024). These advances in "big data biology" include innovations in genetic sequencing, mass spectrometry, bioinformatics, and the development of data-rich biobanks (Gomes and Ashley 2023). In fact, the title of a later report from the National Academies of Science, Engineering, and Medicine ([NASEM] 2016; formerly the NRC) declares that biomarker tests are "key to unlocking precision medicine." Thus I will focus on the narrow form in the rest of this

section. I will briefly discuss (narrow) radical personalized medicine before turning to (narrow) precision medicine.

Narrow radical personalized medicine involves making predictions about an individual patient based on their unique molecular biology rather than the type of disease they have. The strategies proposed for doing so are diverse and include pathophysiological reasoning, computer modeling, and patient-derived organoids.[1] Pathophysiological reasoning involves making inferences about a patient's disease or response to medical intervention from an understanding of relevant molecular mechanisms combined with molecular data from the patient (Lange 2007; Tonelli and Shirts 2017; Montano et al. 2022). Computer modeling could involve simulating mechanisms (computer-assisted pathophysiological reasoning) as well as simulating protein structure and function using information about a patient's genotype (Walker et al. 2019; Tonelli and Shirts 2017). Finally, patient-derived organoids (or organs-on-a-chip), such as those produced using induced pluripotent stem cells created from a patient's own cells, are tissue culture models that can be used to recapitulate a patient's genotypically unique tumor (for instance) for the purpose of screening drugs in vitro that can then be used in the same patient (Letai 2017; Walker et al. 2019; Bose et al. 2021; Green et al. 2022).

To the extent that these models avoid inferences from epidemiologic studies and involve making inferences directly from theory, computer modeling, or model systems, they provide a stark contrast to epidemiological medicine and represent a new kind of personalized, molecular laboratory medicine. However, while these strategies are common in preclinical research to aid in drug discovery and development, they are less common as preferred models for scientific medicine in which their results are directly integrated into clinical care. Pathophysiological reasoning involving molecular patient data such as whole genome or exome sequencing is sometimes used in the state-of-the-art diagnosis and treatment of patients with rare or novel genetic diseases for whom epidemiological data are scant or nonexistent, particularly at specialized centers and as part of research (Montano et al. 2022). Similarly, patient-derived organoids and other laboratory studies have sometimes been used to test the effectiveness of treatments for patients

[1] Examples of the broad form of radical personalized medicine include *n*-of-1 clinical trials and 3D printing to generate prostheses guided by medical imaging of the particular person's anatomy.

with rare mutations, bypassing clinical trials (Tonelli and Shirts 2017; Bose et al. 2021). Yet these contexts are exceptional.

Moreover, some visions for the way these results will be integrated into clinical medicine may be more similar to epidemiological medicine than first meets the eye. For instance, it is generally agreed that patient-derived organoids (PDOs) must typically be validated in clinical studies before they can be deployed in clinical practice. This has often involved measuring the predictive value of results in a PDO drug screen for similar results in a patient trial, and one vision for PDO research uses PDO response to stratify patients into subgroups in a clinical trial or observational epidemiological study to measure subgroup-specific treatment effects (Bose et al. 2021). Clinical predictions can then be made directly from the epidemiological research results. This model for "PDO personalized medicine" sees PDO response as a mere biomarker akin to other laboratory assays currently used in epidemiological medicine. (As we will see, this is similar to the role of molecular biomarkers in precision medicine.)[2] Thus, while models of truly radical personalized medicine are in development and are sometimes deployed, not all models involving these new technologies are such a radical departure from epidemiological medicine. And those that are radical are less widely funded and supported in research than is precision medicine. Thus I will focus on the more conservative vision of narrow precision medicine, which is a stronger candidate for the near-term future of scientific medicine.

In "Toward Precision Medicine," the NRC used a definition of precision medicine borrowed from the President's Council of Advisors on Science and Technology (2008, 1): "the tailoring of medical treatment to the individual characteristics of each patient . . . to classify individuals into subpopulations that differ in their susceptibility to a particular disease or their response to a specific treatment. Preventative or therapeutic interventions can then be concentrated on those who will benefit, sparing expense and side effects for those who will not." According to the report, these characteristics include genomic, epigenomic, metabolomic, and exposure data.[3] Precision medicine

[2] Mary Walker and colleagues (2019) suggest that the correlational (epidemiological) studies are used differently in PDO validation than they are in traditional (epidemiological) medicine and in what I am calling "precision medicine." According to Walker and colleagues, in PDO validation, correlational studies validate the inference from the PDO to the patient, while in these other approaches the inference is from the correlational study results to the patient. However, insofar as PDOs are validated for specific diseases and treatments and their predictive accuracy or the PDO response-stratified treatment effect is quantified (Bose et al. 2021), PDO response is treated much like the results of a biomarker test in traditional medicine and precision medicine.

[3] Sebastian Schleidgen and colleagues provide a similar definition of personalized or precision medicine (PM), a precising definition based on a systematic review of the medical literature: "PM

is partly a response to the problem of heterogeneity: patient groups vary in their risk of disease and in their response to medical intervention, yet average or aggregate measures in clinical populations classified by current disease categories mask this heterogeneity. The precision medicine solution is to stratify patients in a way that minimizes this heterogeneity by using molecular diagnostics to recategorize patients and then selectively treating patients with particular molecular lesions, often using a targeted molecular intervention. A common synonym for precision medicine is thus "stratified medicine," and the ultimate long-term aim of the 2011 NRC report is a new taxonomy of disease. New disease categories could be molecular subtypes of existing categories or new types of disease that pave over existing categories.

Precision treatment usually has two components: a biomarker test to stratify patients by molecular type, and a treatment given to patients according to their molecular type (FDA 2013). An exemplar commonly used in the precision medicine literature to illustrate this strategy (though one preceding the precision medicine movement itself) is treatment for HER2-positive breast cancer. Tumors in which the HER2 cell surface receptor is overexpressed have a worse prognosis. In the 1990s, a molecular diagnostic was approved by the U.S. Food and Drug Administration (FDA) for diagnosing people with HER2 overexpression, and a monoclonal antibody called trastuzumab was approved for treatment of breast tumors with HER2 overexpression. As a monoclonal antibody against HER2, trastuzumab is a "targeted therapy" that selectively binds to HER2 and treats tumors with HER2 overexpression.[4] "HER2-positive breast cancer" thus became a standard molecular subtype of breast cancer. Because monoclonal antibodies (mabs), with their ability to bind selectively to protein variants, have been such an important paradigm for precision medicine, Maël Lemoine (2017b) calls precision medicine "mab medicine." And because genetic and epigenetic changes play such an important role in the pathophysiology of cancer (according to the somatic mutation theory discussed in Chapter 2), precision oncology is the prototype precision medicine specialty (Plutynski 2022).

seeks to improve stratification and timing of health care by utilizing biological information and biomarkers on the level of molecular disease pathways, genetics, proteomics as well as metabolomics" (2013, 10).

[4] While targeted molecular therapies are the most characteristic treatment modalities in precision medicine, non-targeted therapies are also used. Another precision medicine exemplar commonly cited is the diagnosis of BRCA mutations that mark patients at higher risk of breast cancer, for whom bilateral mastectomy will provide a greater absolute reduction in risk of breast cancer (NRC 2011a).

Chin-Yee (2024) argues that precision medicine is not merely biomarker-stratified medicine and that targeted treatment is central to the precision medicine paradigm. He focuses on precision oncology, but his point can be generalized to the rest of precision medicine. According to Chin-Yee, precision medicine is a paradigm for treatment, one in which the treatment selectively targets the molecular biomarker in question the way that trastuzumab selectively binds to HER2. This characterization is consistent with many exemplars of precision medicine, including trastuzumab, pertuzumab, imatinib, gefitinib, and pembrolizumab (Chin-Yee 2024). It also accords with the focus on biomarkers for molecularly targeted therapies in NASEM's 2016 report. On the other hand, it rules out some precision medicine exemplars, including in prognosis, and is more restrictive than the definition used in the landmark 2011 NRC Report. However, little hangs on whether I adopt this more restrictive understanding of precision medicine. The main point I now want to make is about the kind of science that grounds precision medicine, which applies well to both Chin-Yee's more restrictive characterization and the NRC's (2011a) slightly more expansive one.

The precision medicine movement is more than just a strategy for tailoring diagnosis and treatment to individuals. It also includes strategies for research that can help achieve these aims. While precision medicine research is diverse, a major priority of the precision medicine community has been data-driven biological research (including big data) as well as innovations in clinical trials.[5] In particular, precision medicine research includes genome-wide association studies (GWAS), -omics-centered cohort studies, and biomarker-based clinical trials. GWAS, a methodology increasingly in use since the 1990s, aim to identify associations between disease or clinical traits and genomic variability, most commonly single nucleotide polymorphisms. They typically have a case-control design (Uffelmann et al. 2021). Such studies are common within the NIH's Cancer Genome Atlas project, which analyzes samples from thousands of individuals with dozens of cancer types using whole genome sequencing and other -omics screening methods (Hutter and Zenklusen 2018). -Omics-centered cohort studies such as those taking place under the NIH's signature All of Us Precision Medicine Initiative (Denny et al. 2019) follow patients with molecular variability forward in time to correlate their molecular features with future disease or

[5] NRC (2011a), FDA (2013), F. Collins and Varmus (2015), Andreoletti (2018), Paneth and Joyner (2018), Green et al. (2019), Phillips (2020).

clinical outcomes. All of Us is an ongoing program that seeks to enroll one million Americans over time and collect from them biospecimens, electronic health record data, and survey data that can be analyzed in many studies that draw on this common data pool.

Finally, biomarker-based clinical trials analyze treatment outcomes within groups classified by molecular biomarkers. In precision oncology, biomarker-based trials use a variety of approaches adapted from traditional clinical trial designs (J. Park et al. 2020; Chin-Yee 2024), including enrolling participants with a single disease and biomarker ("enrichment"), enrolling participants with a single disease and stratifying them according to different biomarkers ("umbrella trials"), and enrolling participants with a single biomarker who have different diseases ("basket trials"). An example of a biomarker-based trial is the landmark I-SPY 2 trial (H. Wang and Yee 2019), an umbrella trial for patients with locally invasive breast cancer that has studied precision oncology drugs like pertuzumab and pembrolizumab, which are viewed as precision medicine exemplars.

What's notable about these precision medicine studies is that they are essentially epidemiological. The origins of precision medicine are not just in laboratory science but also in epidemiology and biostatistics (Phillips 2020). GWAS, -omics cohort studies, and biomarker-based trials are specially designed to handle the feast of having large amounts of molecular data on the one hand, and the famine of small clinical trial samples consisting of patients with rare molecular signatures on the other. Nonetheless, they are still derived from classical epidemiological study designs: the case-control study, the cohort study, and the clinical trial.

As a variety of scientific medicine, precision medicine is (or is to be) integrated with the results of these studies in several ways. First, as these are clinical studies, research and practice are necessarily interwoven. Second, these studies can yield results relevant to clinical medicine, including discovery of the molecular causes of diseases, molecular risk prediction algorithms, identification of biomarkers that can be exploited in diagnosis and rational drug design, and treatment effect estimates stratified by molecular biomarkers. Precision medicine thus resembles classical epidemiological medicine in being integrated with epidemiological research (with narrow precision medicine based more specifically in *molecular* epidemiology).

Although it is still early days, there are signs that precision medicine is also being modeled after epidemiological science in ways that resemble classic epidemiological medicine. While the emphasis on molecular pathology and

genetics has increased with precision medicine, the emphasis is not brand new. Disease concepts and theory have embraced molecular explanations since the second half of the 1900s, especially in the areas of cancer and Mendelian genetic diseases.[6] Narrow precision medicine's intense focus on genes represents a difference in degree rather than kind compared to classical epidemiological medicine.

Moreover, ever since the results of early GWAS found that genetic features accounted for only a small amount of variation in the incidence of common diseases (L. Ross 2023), GWAS and -omics research more generally have embraced the concept and theory of multifactorial etiology. For instance, one output of contemporary GWAS are polygenic risk scores (Uffelmann et al. 2021), which can consider multiple genetic polymorphisms within multivariate risk prediction algorithms in the way that Framingham initially did for cardiovascular risk factors. In fact, Framingham itself now incorporates genomic analysis (Govindaraju et al. 2008). A longer-term aim of precision medicine research is to include non-genomic variables, including other molecular data, exposure data, and social/behavioral data (NRC 2011a), in order to truly get a handle on the multifactorial causation of contemporary diseases. While the genetic polymorphisms found in GWAS might not themselves cause the traits they predict, the predictive success of polygenic risk scores that include multiple polymorphisms supports a multifactorial picture of genetic causation (Woodward and Kendler 2023), much like Framingham supported a multifactorial picture of the etiology of heart disease despite its risk factors being primarily prognostic (Giroux 2012; Aronowitz 2015).

Also evident in the use of GWAS and polygenic risk scores is the persistence of risk prediction and risk reduction as important goals in precision medicine. A common output of these methods is individualized risk estimates. Moreover, one component of the precision medicine strategy is to stratify patients by biomarkers that better sort higher-risk from lower-risk clinical populations so that interventions from screening to surgery can be

[6] The strategy of molecular diagnosis, described as identifying the "intrinsic biology" of disease (NRC 2011a), has even more distant origins in the principle of classifying diseases according to their specific pathology (Chapter 3), albeit with a new focus on *molecular* pathology. Similarly, the idea of targeted therapy, treatments that target disease pathology with high specificity by binding to it selectively the way that a monoclonal antibody binds to a specific antigen, is the old concept of a magic bullet coined by the chemist Paul Ehrlich at the beginning of the 1900s in thinking about chemotherapies that could selectively target germs. Foreshadowing the use of monoclonal antibodies in oncology, Ehrlich's concept of a magic bullet was derived from his primordial theory of antibodies, and he even proposed extending the magic bullet strategy to cancer therapy (Mukerjee 2010).

offered selectively to the higher-risk populations. Again, this strategy is not new, merely modified by the preference for molecular biomarkers in doing the stratifying. One common strain of precision medicine called P4 medicine even aspires to be *predictive* and *preventive* (Green and Vogt 2016), explicitly adopting the aims of the new preventive medicine of the second half of the 1900s (discussed in Chapter 6): risk prediction and risk reduction.

Finally, the epistemology of precision medicine is similar to classic epidemiological medicine in that both models are guided directly by epidemiological evidence and research results. In precision medicine, the results of GWAS and biomarker-based cohort studies inform prognosis, while the results of biomarker-based clinical trials inform intervention. As in classic epidemiological medicine, clinical trials are generally required to validate findings from discovery research and for FDA drug approval, though drug regulation in precision medicine is evolving (FDA 2013).

Precision medicine trials have sometimes been criticized by EBM traditionalists as failing to live up to high methodologic standards (Prasad 2020). Several recent FDA approvals have been granted for precision therapies (typically in the treatment of advanced cancers) on the basis of small, uncontrolled trials using surrogate endpoints (Tonelli and Shirts 2017; Andreoletti 2018). However, many of these trials that are disparaged by critics or used to support regulatory approval are early-phase trials. Further, the approval in question is often expedited approval, and RCTs using survival endpoints are still commonly used for full drug approval (Beaver et al. 2018). For instance, the I-SPY-2 trial uses a surrogate endpoint, but it is a phase II screening trial meant to speed discovery of promising agents that can then be studied in confirmatory randomized trials (H. Wang and Yee 2019). Therefore it is often better to view precision medicine trials as an innovation in early-phase trial design rather than an alternative to EBM's favored phase III randomized trials. Moreover, it is not a necessary feature of precision medicine trials that they are uncontrolled or use surrogate endpoints. Umbrella and basket trials may even be RCTs (J. Park et al. 2020), as in the example of I-SPY-2.

That said, confirmatory precision medicine trials have also been criticized for utilizing surrogate endpoints and departing from gold standard RCTs in other ways (Gyawali et al. 2023).

Moreover, precision medicine trials often use nonclassical designs that are generally becoming more common in medicine, such as adaptive randomization (Andreoletti 2018), in which the probability of assigning the next

patient to a trial group is updated based on the ongoing trial results. Thus the precision medicine therapeutic epistemology may deviate from the EBM standards of evidence that are typical of classical epidemiological medicine.

Finally, precision medicine's population research sometimes uses machine learning rather than conventional statistical analysis (Quazi 2022).[7] While the implications of the use of machine learning for epidemiological medicine are difficult to anticipate, it is worth noting that both epidemiological research with conventional statistics and the use of machine learning algorithms to analyze epidemiological studies share certain epistemically relevant characteristics; both are population-based, and both are often described as "black box" correlational methods (Susser and Susser 1996; Walker et al. 2019), methods that may not require (and in the case of machine learning, may not allow) deep understanding of the underlying mechanisms generating the correlations uncovered. On the other hand, the integration of theory-rich basic molecular biology with precision medicine research could counterbalance the use of theoretically frugal black box methods.

Overall in current theory and practice, precision medicine resembles classical epidemiological medicine in several ways (with a narrowed molecular focus). Precision medicine's epidemiological research, including research on its targeted therapies and biomarkers, is supported by upstream laboratory research, including the mechanism discovery, computer modeling, and in vitro screening studies discussed earlier as strategies for radical personalized medicine. Yet the same is true of classical epidemiological medicine. What they also share in common is that in both models the epidemiology has greater clinical priority, more directly justifying the clinical use of biomarkers and treatments compared to the results of the laboratory research. One notable difference between classical epidemiological medicine and precision medicine seems to be their therapeutic attitude. Classical epidemiological medicine was inherently skeptical as it tested existing technologies provided by a distinct laboratory culture and provided new standards for their regulation and use. Perhaps because precision medicine is more closely tied to the laboratory research generating new molecular diagnostics and treatments, it seems comparatively more optimistic and susceptible to hype.

[7] Another potential future model of mainstream scientific medicine is "deep medicine" (Topol 2022), in which the science of deep learning or machine learning becomes the characteristic science of scientific medicine. However, machine learning is also a set of methods that can be used in various sciences, just like conventional statistics, which is why I consider it here as a development within epidemiological medicine.

Looking forward, whether precision medicine will eventually become another new modern scientific medicine will depend on the extent to which radical personalized medicine proposals are realized, the effects of machine learning and other nontraditional methods of data analysis on scientific medicine's concepts and epistemology, and the extent to which precision medicine supplements molecular variables with traditional clinical variables. It also depends on whether precision medicine arrives at all beyond oncology and the treatment of genetic disorders. Precision medicine is currently a reality for only a minority of diseases and is an area with perhaps more hype than concrete results (Gamma 2017; Tabery 2023; Plutynski 2022). If precision medicine turns out to be what is essentially a new epidemiological medicine, then it may inherit long-standing problems of the new modern medicine discussed in section 9.2 (if it has not already), such as molecular reductionism, multifactorial causation, the meaning of medical risk, biased evidence, and—ironically for a model concerned with tailoring medical care to the individual—the challenges of making inferences about individuals from population data. Contrariwise, the more precision medicine deviates from classic epidemiological medicine, the more it may raise new problems about, for instance, the reliability of biological theory in guiding clinical reasoning and decision-making.

For now, the leading vision of personalized/precision medicine is not one in which epidemiology is abandoned in favor of laboratory molecular biology in a medical-scientific revolution, but rather of a molecular epidemiological medicine, of a clinical medicine integrated with and modeled after molecular epidemiology. Precision medicine is best viewed as a development within our current model of scientific medicine rather than as a radically new modern scientific medicine.

## 9.4. Postmodern Medicine

While precision medicine promises a revolution within medicine that could bring about yet another new scientific medicine, other recent developments in medicine threaten an anti-scientific medical revolution in society in which scientific medicine would lose its privileged role as mainstream medicine. Pseudoscientific medicine seems to be on the ascent, while scientific medicine seems to be in decline. If we understand modern medicine, including the new modern medicine I've written about in this book,

as defined by the societal dominance of scientific medicine, then scientific medicine's changing fortunes could mark the end of modern medicine, giving way to a postmodern medicine in which scientific and pseudoscientific medicines enjoy equal social status. In this section, I will briefly examine these possibilities, using as case studies homeopathy (section 9.4.1) and the COVID-19 pandemic (section 9.4.2). I conclude that the stark concern that postmodern medicine is right around the corner is exaggerated, although these examples should prompt serious reflection on our current model of scientific medicine.

For decades, the use of complementary and alternative medicine has been increasing around the world (Frass et al. 2012; Nahin et al. 2024). Institutions for research into complementary and alternative medicine have popped up, including the National Center for Complementary and Integrative Health (NCCIH 2021a) of the NIH, founded in 1998. Meanwhile, there are growing pockets of public skepticism about mainstream medical interventions, most notably, hesitancy about vaccines, including childhood vaccines and COVID-19 vaccines (Goldenberg 2021; Hotez 2023). Vaccine skepticism is often seen as a marker of distrust in mainstream scientific medicine. Anti-institutionalist political movements around the world threaten the authority of scientific medical institutions (Hotez 2023). The legitimacy or adequacy of psychiatric diagnosis and psychiatry is sometimes questioned by patient activists such as the neurodiversity and mad pride/mad studies movements (Arnaud and Gagné-Julien 2024).[8] Even within medicine there is often a sense that medicine has passed its golden age (Varga 2022) and a sentiment among some that modern medicine needs rethinking (Marshall et al. 2018). We encountered concerns about the influence of private interests and the effectiveness of medical interventions, including concerns among medical insiders, in Chapter 8. To these concerns we could add mainstream medical worries about overmedicalization, overdiagnosis, and "too much medicine" (Moynihan and Smith 2002; Moynihan and Cassels 2006; Welch et al. 2011; Le Fanu 2019). Overall there might seem to be a mainstreaming of pseudoscientific medicine, coupled with the threat of a diminished status for mainstream scientific medicine.

From this review of scientific medicine's recent health history, we might wonder whether modern medicine is in decline as we enter a new era of

[8] However, Sarah Arnaud and Anne-Marie Gagné-Julien (2024) argue that these contemporary movements have taken a *scientific* turn.

postmodern medicine. Just as postmodernism as an intellectual movement is associated with skepticism about science and claims of its privileged objectivity as well as a pluralism regarding perspectives on the world, postmodern medicine would mark a skepticism about scientific medicine, the undermining of its privileged place in society, and the welcoming of alternatives to mainstream scientific medicine (Gray 1999). Most important, scientific medicine would lose its position of societal dominance, in both the state and wider society. The resulting medical eclecticism could recall the wide-open marketplace of early nineteenth-century American medicine, in which practitioners of scientific medicine competed with less educated orthodox doctors as well as homeopaths and quacks (Warner 1997, 1998). Do recent trends indeed portend this fate for the new modern medicine?

Rather than address this dauntingly complex question in all its dimensions, I will look for clues in two important case studies: homeopathy and COVID-19. Homeopathy is a paradigm exemplar of pseudoscientific medicine and an indicator of the status of pseudoscientific medicine in society, while the COVID-19 pandemic was a crucial test of scientific medicine's authority. I won't discuss all the harbingers of postmodern medicine mentioned above, but these examples are testing grounds for many of them. While both cases pose significant challenges for the new modern medicine, they illustrate as many—if not more—indicators of the power of epidemiological medicine as they do signs of its impending demise.

### 9.4.1 The Rise of Pseudoscientific Medicine: The Case of Homeopathy

From the beginning of the scientific medical tradition in the nineteenth century, homeopathy has been used as a pseudoscientific contrast case to help define a scientific medical identity. Classical homeopathy persists today. Some of its recent research, especially systematic reviews of RCTs, appears indistinguishable from scientific medicine's research and seems to support the homeopathic system. Assuming that homeopathy is still unscientific, this development challenges the scientific credentials of epidemiological medicine and its exclusive grip on healthcare. Yet I will argue that scientific medicine can continue to demarcate itself from unscientific homeopathy today by modifying a "like comparison criterion" that it originally used in the nineteenth century against earlier homeopaths.

In Chapter 1, I discussed what was perhaps the first popular model of modern scientific medicine, "hospital medicine" (Ackerknecht 1967; Jewson 1976) or "Paris clinical medicine." This model was imported into America from Paris in the early 1800s. Parisian doctors teaching in large hospitals directed the revolutionary spirit of the era at clinical medicine. I say that theirs was a model of modern scientific medicine because they integrated what we would today recognize as sciences into clinical care, in particular the clinical sciences of pathological anatomy and clinical statistics. The latter involved the "numerical method" of P. C. A. Louis, which required classifying cases, counting outcomes, and—crucially—comparing groups. In therapeutics, Louis described and often used what I call a "like comparison criterion" for evaluating treatments for clinical use. I give the criterion this name because it uses "like comparisons," comparisons in which we are comparing "like with like" (Fuller 2024). Scientific therapeutic research involved comparing cases that "were in all respects alike, except in the treatment," to see which treatments were associated with the greatest recovery (Louis 1837, 490–491). Scientific therapeutic practice then integrated the results of these like comparisons to inform treatment. Louis didn't invent this idea, but he did make it part of the foundation for his model of medicine.

Foreign doctors who trained under the Paris clinicians brought the model back home with them, often settling into jobs in universities. American doctors were returning home to a highly eclectic medical landscape. They were in competition with unscientific practitioners of all kinds, including orthodox doctors who did not receive the same scientific training but who were raised on a similar diet of conventional therapies like bloodletting, as well as heterodox doctors who subscribed to alternative therapeutic systems (Warner 1997, 1998). Among the heterodox medical sects, homeopathy was the most popular, particularly among the wealthy class.

Not only were the homeopaths serious marketplace competition for the American purveyors of Paris clinical medicine, but the homeopathic system offended their scientific sensibilities. Classical homeopathy was a system developed by the German physician Samuel Hahnemann (1810) in the late 1700s and early 1800s. Its two enduring principles are the law of similars and the principle of dilution. The law of similars holds that "like cures like": one can cure a disease characterized by certain symptoms using a substance that produces those symptoms in people without the disease. The principle of dilution holds that one should greatly dilute that substance before administering it through serial dilution, shaking the mixture at each dilution

stage. Homeopathic treatments are typically highly diluted, often so much so that not a single molecule of the substance is likely to remain.

Practitioners of Paris clinical medicine like Oliver Wendell Holmes, a future dean of the Harvard Medical School, ridiculed Hahnemann and his followers. In *Homeopathy and Its Kindred Delusions*, Holmes implied that homeopathy, like astrology and alchemy, was a "pseudoscience" (1842, 1). Holmes noted that homeopathy's principles were inconsistent with the science of the day. But Holmes also took aim at the evidence that homeopaths used to justify the use of particular homeopathic treatments. Such evidence typically consisted of noncomparative case studies or anecdotes. When symptomatic remission followed homeopathic treatment in these cases, the homeopaths fallaciously attributed recovery to the treatment. Holmes and other Paris disciples argued that the scientific way to assess treatment efficacy was instead through like comparisons. They wielded the like comparison criterion to distinguish their scientific medicine from the unscientific medicine of the homeopaths (Fuller 2025). An example of a like comparison they frequently lauded was Louis's (1836) researches on bloodletting, in which Louis compared the length of illness among hospital inpatients bled earlier in the course of disease with the length of illness among inpatients bled later.

The like comparison criterion served as a good "demarcating difference" between scientific medicine and unscientific medicine when it came to nineteenth-century Paris clinical medicine and homeopathy. It invoked a scientific principle of causal inference—that in inferring causation we should hold other causes fixed—and marked an important difference between Paris clinical medicine and homeopathy as models of medicine. It allowed practitioners of Paris clinical medicine to claim a unique scientific identity for themselves. Eventually, orthodox medicine collectively adopted science as its main source of professional identity, and in the early twentieth century began its rise to dominance as modern scientific medicine (Chapter 1).

However, it might seem that the like comparison criterion can no longer successfully demarcate contemporary scientific medicine, represented by epidemiological medicine, from contemporary unscientific medicine, represented by classical homeopathy. Classical homeopathy persists today and is consumed by not insignificant numbers of people in America and worldwide (Clarke et al. 2015; Relton et al. 2017). One thing that has changed is that there are now RCTs that might seem to justify this use. In particular, several systematic reviews of RCTs with meta-analysis seem to

measure statistically significant pooled effects of a diverse range of homeopathic remedies compared to placebo (Mathie et al. 2014, 2017). While these effects are small, the authors note that so too are the effects of many mainstream medical interventions (Mathie et al. 2014). In fact, these reviews are similar to those conducted in mainstream scientific medicine, done according to EBM standards for systematic review. The results are particularly important because they appear to confirm the principles of classical homeopathy by showing effects of highly diluted treatments compared to chemically identical placebos when given to patients according to the principle that like cures like.

When epidemiological medicine's gold standard scientific method for evaluating therapies confirms homeopathy's therapeutic system, it poses a threat to epidemiological medicine's status as mainstream scientific medicine. If medicine maintains that homeopathy remains unscientific,[9] it becomes difficult to draw a line between homeopathy and mainstream EBM using EBM's own standards that is sharp enough to demarcate epidemiological medicine as decidedly more scientific. Pseudoscientific medicine seems to have risen to the scientific standards of epidemiological medicine, revealing those standards to be weak and unable to distinguish epidemiological medicine as uniquely deserving of the title "scientific medicine."

My response to this concern is that it is revealed to be exaggerated once we consider how epidemiological medicine adapts the concept of a like comparison. At the same time, EBM's standards are indeed weak in at least one respect, and we should modify the like comparison criterion to more definitively demarcate scientific medicine from the unscientific medicine of homeopathy. A modified like comparison criterion would interpret a like comparison as an unbiased comparison, consider the mechanistic implausibility of therapies in evaluating the probability of an unbiased comparison, and require that therapeutic decision-making respect the probability that a therapeutic comparison is unbiased (Fuller 2024). I will take each component in turn.

[9] One could simply bite the bullet and accept that homeopathy is in fact an alternative scientific medicine. But since homeopathy has long been seen as the arch pseudoscientific medicine, and homeopathic principles are seen as inconsistent with much contemporary science, that concession would be difficult to swallow. I am bracketing the distinction between unscientific medicine and pseudoscientific medicine here. While the like comparison criterion distinguishes scientific from unscientific medicine, homeopathy is also widely considered to be pseudoscientific because it appears scientific (Mukerji and Ernst 2022), including in the way that it flashes empirical data of the kind discussed here.

Louis's idea of a like therapeutic comparison was one in which the groups being compared were similar in known, observed causes (except for the treatment). In his famous studies of bloodletting, Louis (1836) observed whether the cases bled earlier resembled the cases bled later in factors such as age and number of bleedings. Much later, in Hill's popular textbook *Principles of Medical Statistics* (Hill and Hill 1991), Hill gave an epidemiologist's diagnosis of the problem with such studies: the way that the groups were selected could be biased, resulting in one group having a worse prognosis. (For instance, he suggested that patients treated later tend to be those who have failed to recover or who are worsening.) Hill's main reason for introducing randomization into the 1948 streptomycin trial was that it eliminated an important source of selection bias, namely, bias that occurred when trial recruiters could predict the group to which a participant would be allocated and decide to enroll them or not on that basis (Chalmers 2005). When the allocation is random, it is impossible to predict, at least until the allocation schedule has been set. Thus, for epidemiologists like Hill, a like comparison is an unbiased comparison.

Epidemiological medicine inherits epidemiology's interpretation of a like comparison. In EBM, a central determinant of the quality of therapeutic evidence is its risk of study bias. As the authors of the homeopathy reviews note, almost none of the RCTs included were at low risk of bias (Mathie et al. 2014, 2017). Thus, according to a proper like comparison criterion that keeps up with the changing science of scientific medicine, the homeopathy RCTs are probably *not* making a like (i.e., unbiased) comparison.

To be fair, many RCTs of mainstream medical interventions are too at a substantial risk of bias (Howick et al. 2022). However, the risk of bias in epidemiology is evaluated mainly according to methodologic features of the study. The probability that a positive study is biased should also be influenced by the probability of alternative explanations for the positive study result that compete with the bias explanation (Howick 2011). These competing explanations include chance (ruled out by classical statistics) and the hypothesis that the treatment has efficacy. Thus, in addition to directly assessing the probability of study bias by appraising study features (as EBM does), we can additionally *indirectly* assess the probability of bias by evaluating the probability of the efficacy hypothesis.

One way to independently evaluate the efficacy hypothesis is through the basic science or mechanistic evidence that EBM spurns but that some philosophers argue should be evaluated alongside the difference-making

evidence provided by RCTs (Clarke et al. 2014; J. Stegenga 2022). However, we don't need to evaluate the degree to which therapeutic mechanisms are plausible in order to more sharply demarcate homeopathy from scientific medicine. It's commonly argued that it is *categorically implausible* that homeopathic treatments have effects beyond placebo (Turner 2017). If that is right, then a modified like comparison criterion need consider only categorical mechanistic implausibility in assessing the probability that bias is the best explanation for a study's result. With this modification, homeopathy RCTs are much more likely to be biased compared to RCTs of mainstream medical interventions because while proposed mechanisms of homeopathic treatment are implausible, proposed mechanisms of mainstream treatment are generally not.

Finally, a modified like comparison criterion should require that scientific practice respect the probability that a therapeutic comparison is unbiased (while letting mechanistic implausibility influence this probability). In EBM's popular GRADE approach to evaluating evidence for decision-making, risk of bias is an important determinant of quality of evidence,[10] and quality of evidence is an important determinant of strength of recommendations for/against using treatments (Balshem et al. 2011). In contrast, classical homeopathy lacks high-quality, unbiased therapeutic comparisons yet remains stubbornly confident in the efficacy of its therapeutic system (Ernst 2016). Thus, by the modified like comparison criterion, epidemiological medicine is scientific, and classical homeopathy is comparatively unscientific. If epidemiological medicine is willing to allow judgments of mechanistic implausibility to complement epidemiological risk-of-bias assessments, it will have a powerful criterion for defending its scientific credentials against the challenge from homeopathic pseudoscience.

By and large, while global consumption of homeopathy has remained steady in recent decades, it is not increasing (Clarke et al. 2015; Relton et al. 2017). Meanwhile, in some places, state support (or better put, state tolerance) has decreased (Ullman 2021). In 2009, the U.K. House of Commons Science and Technology Committee concluded from the available RCT evidence that homeopathy is not efficacious (Turner 2017); in 2017, the National Health Service in England announced it would no longer be

[10] Though we saw in Chapter 8 that not all kinds of bias are recognized in GRADE. In particular, GRADE does not formally consider industry-sponsorship bias or bias arising from conflicts of interest.

funding homeopathy prescriptions. In 2015, Australia's National Health and Medical Research Council agreed with the U.K.'s assessment of the evidence on homeopathy. The NIH's center for complementary medicine, the NCCIH (2021b), reiterates this consensus.[11] The U.S. FDA regulates homeopathic drugs as it does any other drug, requiring epidemiological evidence of safety and efficacy (though enforcement is targeted and selective). No homeopathic drugs are currently approved by the FDA (2022). Therefore homeopathy still does not have the state support or wider popularity that scientific medicine enjoys.

As long as there has been a mainstream scientific establishment, there has also been a fringe counter-establishment (Gordin 2021). But at least when it comes to homeopathy, the fringe has not yet become mainstream. Homeopathy is not on the equal footing with orthodox medicine that it enjoyed in the 1800s, before the age of modern scientific medicine; medicine is not experiencing a postmodern revival of that premodern state of affairs. While the recent epidemiological research discussed could lend some scientific legitimacy to homeopathy, epidemiological medicine could also strengthen its own scientific credentials and ward off this challenge from homeopathic pseudoscience simply by modifying its standards for making like comparisons to include assessments of mechanistic implausibility. It could also respect the many potential sources of bias in contemporary RCTs, including biases detected by meta-research evidence (discussed in Chapter 8).

### 9.4.2 The Fall of Scientific Medicine: The Case of COVID-19

In section 9.4.1, I discussed one way in which scientific medicine could lose its privileged place in society, bringing about a postmodern medicine: pseudoscientific medicine could ascend to scientific medicine's pedestal. Another route to postmodern medicine would be the decline of scientific medicine, in

[11] Aside from homeopathy, the use of complementary and alternative medicines has been increasing for decades (Frass et al. 2012; Nahin et al. 2024). One might worry that the emergence of institutes like NCCIH provides mainstream legitimacy to pseudoscientific medicine. However, NCCIH's mission is to scientifically evaluate complementary medicines, which it defines as medicines used alongside conventional therapies (as opposed to alternative medicines, which are used instead of conventional therapies; NCCIH 2021a). Complementary medicines that are demonstrated to be effective in scientific research and are used within mainstream healthcare alongside conventional treatments arguably become part of mainstream scientific medicine. Rather than undermining mainstream scientific medicine's authority and reach, centers like NCCIH seek to expand it.

which it slips from its position of power. In this subsection, I look for further clues into modern medicine's future from a second case study, the COVID-19 pandemic. We find signs of resistance to scientific medicine, but also signs of its strength and, just as important, signs of its limitations in the midst of a formidable twenty-first-century crisis.

This book began in 1874 with the story of Dr. Osler attending to a man dying from smallpox in Montreal. A decade later, the city endured North America's worst ever smallpox epidemic, despite the availability of smallpox vaccines. The reasons why Montreal was so susceptible to an outbreak of smallpox compared to other cities are varied but include a skepticism among a large portion of the population at the time toward the secular, scientific solution to the epidemic provided by smallpox vaccination (Bliss 2011). Smallpox was once one of the unstoppable pandemics of premodern medicine, a period in which infectious diseases were the leading cause of death in Western cities like Montreal, and scientific theories and solutions that could effectively combat them were scarce. By the time smallpox was eradicated by a worldwide vaccination campaign in the 1980s, scientific medicine had become a mainstream global institution. Modern medicine had arrived.

Yet despite premature claims that Western medicine had turned the page on infectious disease pandemics, around the time that medicine was declaring victory over smallpox the threat of emerging and re-emerging infectious diseases was gradually becoming recognized (Snowden 2008). Emerging infectious diseases are epidemic infections like HIV disease, Ebola virus disease, and COVID-19 that are due to novel pathogens, while re-emerging diseases are infections like tuberculosis (including multi-drug-resistant tuberculosis) due to older pathogens with increasing incidence or geographic range (Snowden 2008). Although emerging diseases have occurred throughout recorded history (Morens et al. 2008), the number of emerging and re-emerging infectious diseases seems to be growing (K. Jones et al. 2008). They have proven more intractable in developed countries than one might have expected given the size of medicine's scientific arsenal. Moreover, there is a recalcitrant and rising vaccine skepticism among some members of the public, including hesitancy toward childhood vaccines that for decades had helped to suppress infectious disease outbreaks (Larson et al. 2011; Wiegand et al. 2023).

The COVID-19 pandemic that began in 2020 was an important test for scientific medicine, a crisis that strained its authority and principles, exposing both strengths but also fissures. COVID-19 (coronavirus disease

2019) is a respiratory illness caused by a novel human coronavirus, SARS-CoV-2 (severe acute respiratory syndrome coronavirus 2). From its discovery in China in December 2019, the virus quickly spread to the rest of the world. The WHO officially declared a pandemic in March 2020 and declared the public health emergency caused by COVID-19 over in May 2023. COVID-19 proved to be one of the deadliest and most consequential pandemics in the history of modern medicine, engendering comparisons with the 1918 flu pandemic and even premodern medical plagues (Desai 2020). Just as swiftly as it spread around the world, scientific medicine and public health moved to study COVID-19, to develop interventions to fight it, and to implement medical and public health measures to curtail the pandemic. Doctors quickly repurposed existing interventions like drugs, steroids, antivirals, and intensive care to manage severe disease. Scientists then developed new vaccines, monoclonal antibodies, and antivirals that were eventually used by billions.

A notable feature of the pandemic was highly visible public resistance to COVID-19 interventions and science and, plausibly, to the mainstream scientific medicine supporting them. While there was also widespread resistance to public health measures such as masking and policies restricting people's mobility (e.g., stay-at-home orders), I will mostly examine medical interventions because my focus is on medicine rather than public health. There is no universal distinction between a medical intervention and a public health intervention. One principled way of drawing a distinction is to refer to interventions prescribed, dispensed, or implemented by doctors and other clinicians as medical interventions, and interventions recommended or implemented in public health, including by public health agencies and governments, as public health interventions. Because public health agencies make recommendations about drugs and biologics and because clinicians are often viewed as part of the public health workforce, medical interventions are often also public health interventions (which, I will argue, could help explain some of the resistance to them). On the other hand, public health policies or recommendations such as those concerning social distancing and wearing face coverings in public are less naturally viewed as medical interventions.

While one could multiply examples, I will summarize only some of the sources of resistance, focusing on the American context. (There were many important international differences, especially between countries of the Global North and countries of the Global South.) First, some challenged the principles of EBM, including the universal authority of conventional

randomized trials (Greenhalgh 2020a; Paneth et al. 2022). Early in the pandemic, some doctors used treatments off-label for COVID-19 in the absence of RCT evidence, sometimes instead of recruiting patients for ongoing RCTs, even RCTs of those very treatments, believing that their clinical experience with using those treatments provided sufficient rationale (Dominus 2020). Many early studies of existing therapeutics that were candidates for repurposing to treat COVID-19 were small, non-randomized, uncontrolled, or deviated from EBM standards in other ways (Glasziou et al. 2020; Jung et al. 2021).

Moreover, certain drugs such as hydroxychloroquine and ivermectin were promoted by contrarian voices as treatments for COVID-19 and used by not insignificant numbers of people long after clinical trials showed them to be ineffective and mainstream medicine had discarded them (Horby et al. 2020; Abdool Karim and Devnarain 2022; Lazarus et al. 2023). Though these drugs are approved in mainstream medicine for the treatment of other conditions, they were at this point in the pandemic considered non-mainstream medicines for the treatment of COVID-19; thus they functioned as alternative medicines in this context. Because their use did not respect the negative results of higher quality RCTs but rather the results of non-randomized trials or RCTs of lower quality (Glasziou et al. 2020; Karim and Devnarain 2022), the prescribing of these interventions with a high level of confidence in their efficacy could be seen as violating the modified like comparison criterion introduced in section 9.4.1 and thus deemed unscientific (though, strictly speaking, the criterion is meant to distinguish scientific from unscientific models of medicine rather than individual medical practices).

Contrariwise, COVID-19 vaccines that were enthusiastically promoted by mainstream scientific medicine were seemingly rejected by many members of the public who were hesitant about using them (Lazarus et al. 2023). Even if a minority of individuals outright refused vaccination, hesitancy about vaccines that were developed, tested, and endorsed by the scientific establishment could be seen as a harbinger of a smoldering mistrust or rejection of scientific medicine among the public. New COVID-19 antivirals that were found highly effective in people at high risk from severe COVID-19 disease were also sometimes underutilized (Lazarus et al. 2023).

Some vaccine skepticism, along with other non-mainstream views about medical interventions, was fueled by an "infodemic" of medical misinformation promoted by purveyors of conspiracy theories or science denialists (Zarocostas 2020), categories related to or even continuous with

pseudoscience (Gordin 2021). While typically a fringe community, these groups were able to amplify their views during the pandemic through social media and even mainstream media and politics (Hotez 2023), potentially influencing public belief and behavior (Pierri et al. 2022). Misinformation disrespects scientific evidence but can appear scientific (i.e., it is pseudoscientific) when it cites fabricated data or misleading information.

There are certainly signs of potential deterioration of support for mainstream scientific medicine underlying these trends, which could portend a poor or uncertain prognosis for its survival. However, we must put them in perspective. The COVID-19 pandemic was a public health crisis. By its nature as a crisis on the one hand and a public health problem on the other, the pandemic might paint an overly pessimistic picture of the state of scientific medicine. First, as a crisis, an unstable set of circumstances, certain attitudes during the pandemic might not accurately represent attitudes toward scientific medicine generally, or at least might be exaggerated. Reticence to conduct orthodox randomized trials, the gold standard in EBM, might reflect research exceptionalism (London and Kimmelman 2020), the stance that certain situations such as a health crisis might warrant relaxing the typical methodologic standards of medical research for the sake of expedience or humanitarian need. Similar reservations about running orthodox randomized trials surfaced in the West African Ebola virus outbreak of 2014, an earlier health crisis (Adebamowo et al. 2014; Caplan et al. 2015). Moreover, while there were many methodologically weak clinical trials conducted early in the COVID-19 pandemic, there were also many methodologically strong and often innovative randomized trials (WHO 2021), so the methods of EBM were not completely abandoned.

Second, a pandemic is a public health problem and not only a medical problem, and many medical interventions in a pandemic are also considered public health interventions—most notably, vaccines. While medicine is certainly not apolitical, there are nonetheless ways in which public health is closer to politics. For instance, public health concerns the health of communities, and its interventions include policies and other government actions. The highly politicized context of the COVID-19 public health crisis created conditions that polarized beliefs about public health interventions, often along partisan political lines, even among physicians. This was dramatically illustrated by beliefs about the effectiveness/safety of COVID-19 vaccines as well as hydroxychloroquine and ivermectin for COVID-19, which tracked polarization in their prescribing (M. Barnett et al. 2022; Cox

et al. 2023; Levin et al. 2023). In fact, criticisms of the science supporting public health interventions can be proxies for criticisms of the policies and politicians supporting those interventions (Fuller 2020a).[12] When beliefs about the effectiveness of an intervention become polarized in a context in which mainstream scientific medicine comes down clearly on the side of effectiveness or ineffectiveness, then one of the polarized belief factions necessarily winds up endorsing beliefs inconsistent with scientific medicine. While potentially operative in non-pandemic times, the highly politicized climate of a public health crisis may exacerbate these belief dynamics (O'Connor and Weatherall 2020).

Those caveats aside, the public's rejection of mainstream medical beliefs and interventions and its endorsement of non-mainstream beliefs and interventions suggest a concerning diminishing trust in scientific medicine, at least among some groups in some countries.[13] Americans' trust in scientists declined over the pandemic, with 86% reporting having at least a fair amount of confidence in scientists just before the pandemic, compared to 73% in 2023. (Confidence in *medical* scientists declined from 87% to 77%; Pew Research Center 2023). Even within America, trust is unevenly distributed. In 2023, Democrats had greater trust in scientists than did Republicans, with further polarization according to education, religion, and race (Pew Research Center 2023; Cox et al. 2023). This polarization and weakening of trust in science among certain groups likely reflect a problem of trust in institutions generally (Mills 2023), which includes scientific medical institutions like academic medicine, the medical drug and device industry, and government agencies such as the NIH and the FDA. Americans with lower levels of institutional trust express greater vaccine hesitancy and had lower COVID-19 vaccine uptake (Cox et al. 2023).[14] It is important to note that most Americans do still express some trust in scientific medicine and that selective rejection of mainstream scientific advice about COVID-19

[12] Endorsement or rejection of interventions may even be a marker of identity; in this case, political identity. By analogy, endorsement or rejection of bloodletting was a marker of professional identity for orthodox and heterodox medicine, respectively, throughout most of the 1800s (Warner 1997).

[13] The phenomenon is not universal. In a survey of seventeen countries, the proportion of respondents trusting science and scientists in 2023 sat at 88% and 86%, respectively, which was slightly higher than the relevant proportions immediately preceding the pandemic (3M 2022).

[14] Further, Americans—especially Republicans—became significantly less trusting of the FDA and the U.S. Centers for Disease Control and Prevention as reliable sources of vaccine information between 2020 and 2023. Yet 82% still trusted their own doctor or health provider at least "a fair amount" for reliable vaccine information (KFF COVID-19 Vaccine Monitor 2023), indicating more pronounced decay and polarization of trust in faceless scientific medical institutions compared to human practitioners of scientific medicine.

and COVID-19 interventions cannot be generalized to the whole of scientific medicine. However, given indications of a fracturing trust in scientific medicine in some populations, it is reasonable to frame this issue as a crisis of trust (Goldenberg 2021), a potential tipping point for modern scientific medicine, and to look for ways to strengthen both trust in scientific medicine and its trustworthiness in order to prevent scientific medicine's postmodern demise, however likely or unlikely that result may be.

Contrariwise, there were also many signs during the pandemic of epidemiological medicine's robust health. A staggering amount of epidemiological research was done during the pandemic, and much of it continues. This research includes studies modeling the community spread of COVID-19 using various epidemiological modeling approaches (Adam 2020). It also includes observational studies of the natural history of COVID-19, including on its mode of transmission, the emergence of new viral variants, disease spread in healthcare and other settings, and clinical outcomes in various populations, including post-COVID-19 sequelae. These studies took familiar forms, including case studies of local outbreaks, surveillance studies, case-control studies, and cohort studies, which collectively informed medical understanding of the disease, its spread, and its evolution. For example, the NIH's RECOVER initiative is a large research program devoted to the epidemiology of long COVID, a chronic syndrome that can develop after acute COVID-19 infection. Analogous to the NIH's Precision Medicine Initiative, RECOVER (2024) consists of a cohort of thousands of people enrolled in many studies, including studies of electronic health records, clinical trials, and longitudinal cohort studies that seek to identify long COVID risk factors, subtypes, outcomes, and treatments.

There were also many studies on the prevention and treatment of acute COVID-19, including RCTs. For instance, the University of Oxford's RECOVERY trial began at the start of the pandemic as a randomized, open-label platform trial to identify treatments for patients hospitalized with COVID-19. Early on, it evaluated repurposed treatments such as hydroxychloroquine and steroids, delivering influential results (Horby et al. 2020, 2021). Drug companies sponsored their own RCTs of newly designed COVID-19 vaccines, including mRNA vaccines (Polack et al. 2020; Baden et al. 2021). "Natural experiments" or observational epidemiologic studies further analyzed vaccine effectiveness "in the wild" using data systems that included vaccine registries as well as databases like OpenSAFELY, a database drawing patient records from the U.K. NHS (Curtis et al. 2022;

Leonelli 2024). Vaccine effectiveness studies also included cohort studies of clinicians in healthcare systems, who were routinely tested for COVID-19 (Teerawattananon et al. 2022). The results of pandemic studies were integrated into public health recommendations and policy. They were also integrated into clinical medicine in ways typical of epidemiological medicine, their results informing diagnosis, prognosis, treatment, and prevention of COVID-19, as well as individual health behavior.

All the familiar elements of the epidemiological medicine model were on full display during the pandemic. Medicine provided molecular explanations for COVID-19 and its management, often revolving around SARS-CoV-2's spike protein surface receptor, which served as a target for monoclonal antibodies and vaccines. Mutations in the gene encoding the receptor explained the evolution of viral variants, escape from prior immunity, and treatment resistance (Mushebenge et al. 2023). The causes of infection and severe disease were thought to be multifactorial, involving viral factors (infectivity, virulence), immune factors (natural or vaccine immunity, immune-compromising conditions), and comorbidities. Long COVID emerged as a chronic post-COVID disease with unknown and potentially variable pathology (RECOVER 2024). Epidemiological evidence served as the evidence base for medical decision-making, and often countries extrapolated results from their neighbors to predict the course of the pandemic locally. People were sorted into high-risk (e.g., older) and low-risk (e.g., younger) groups, particularly for the purpose of risk-stratified priority-setting as well as recommendations and decision-making (NASEM 2020; Bhimraj et al. 2023). And medical skepticism abounded about unproven COVID-19 treatments, public health measures, and COVID data, stemming from concerns about biased or missing evidence (Ioannidis 2020).

If anything, epidemiological medicine expanded its reach during the COVID-19 pandemic, influencing countless facets of everyday life as well as public health. In the matter of public health, it sometimes overextended itself. While scientific medicine and scientific public health share a contemporary basis in epidemiology, there are limitations to the epidemiological medicine model when applied to a pandemic. I can only gesture to these limitations here. The first three limitations stem from its individualism as a model of patient care rather than a model of population health. First, its multifactorial model of causation is oriented toward the causes of individual cases of disease (Chapter 3), but this perspective gets us only so far in explaining and intervening in epidemics, which require an additional population

perspective that reflects their population dynamics (Fuller 2022a). Further, while the molecular and biological causes of disease (Chapter 2) were often foregrounded and could be well controlled with biomedical interventions, the behavioral and social causes were comparatively less often emphasized and less well controlled with social, behavioral, and environmental interventions. Moreover, the normally tenuous discussion of an "individual's risk" of disease or severe disease (Chapter 6) is even more tenuous during a pandemic because the risk of an infectious disease is plainly not an isolated property of the individual but depends dramatically on the community prevalence of infection as well as the frequency and behavior of human contacts.

Other limitations result from extending the principles of EBM, which are adapted to the medical context, to a public health context. The usual problem of extrapolating from epidemiological study results in EBM (Chapter 7) is aggravated during a pandemic by constant variation and change, whether in the virus, the population setting, human behavior, or population immunity, all of which can dramatically influence the trajectory of an epidemic and the effectiveness of pandemic interventions, greatly undermining the "external validity" of epidemiological studies. Moreover, EBM's hierarchy of evidence as well as its pronounced skepticism about study results and insistent demand for high-quality evidence arguably serve public health poorly, particularly in a crisis. The hierarchy's privileging of RCTs and meta-analysis of RCTs (Chapter 5) runs up against limitations of RCTs for public health interventions (Frieden 2017) and diverse roles served by diverse evidence in public health (Greenhalgh et al. 2022; Fuller et al. 2024). Further, while bias and quality of evidence are always important concerns in medicine and public health (Chapter 8), we do not need ideal evidence before we can act justifiably (Chapter 7), and setting the epistemic bar for action too high can miss crucial windows to protect the public's health, particularly in a crisis.[15]

Modern epidemiological medicine flexed its muscles during the COVID-19 pandemic, even while some segments of the public grew resistant to it. At the same time, limits of epidemiological medicine, particularly in a pandemic, were exposed. I conclude that the stark concern that postmodern medicine is right around the corner is exaggerated, though it would be hubris to make definite predictions about the longer-term future of modern scientific medicine. Still, the cases of homeopathy and COVID-19 should prompt

[15] For a discussion and exchange about these problems for evidence-based public health published in *Boston Review* early in the COVID-19 pandemic, see Greenhalgh (2020b), Fuller (2020a, 2020c), Lipsitch (2020), Ioannidis (2020). For a post-pandemic revisiting of them, see Schulson (2024).

reflection on the problems with our current model of scientific medicine and how to improve scientific medicine if it is to continue to enjoy and to deserve the status that it has been afforded for the past hundred years. These problems for epidemiological medicine include those discussed in section 9.2 and throughout this book.

## 9.5. Conclusion

Epidemiology continues to permeate all aspects of contemporary scientific medicine. Far from being in decline, precision medicine and the COVID-19 pandemic reveal new directions in which epidemiological medicine is developing, as it continues to function as the dominant model of scientific medicine. There are visions in the scientific research community of a radical personalized medicine waiting on the horizon, which would represent a new molecular laboratory medicine. But a more realistic and conservative model, precision medicine, envisions a revised, molecular-oriented epidemiological medicine rooted in genomics and other disciplines of molecular epidemiology. On the other hand, the persistent popularity of the pseudoscientific medicine homeopathy and the positive results of recent homeopathy RCTs as well as rising resistance to scientific medicine during the recent COVID-19 pandemic threaten to undermine the privileged place of scientific medicine in society. Despite these apparent harbingers of an impending postmodern medicine, epidemiological medicine remains in reasonably good health. However, seven problems of the new modern medicine discussed in this book and shown here to result from the influence of epidemiology on medicine should force a rigorous rethinking of epidemiological medicine if it is to continue to serve as our model of scientific medicine in the twenty-first century.

APPENDIX

# Effect Measures and the Sufficient-Component Cause Model

According to the sufficient-component cause (SCC) model, medical interventions are component causes within sufficient causes of health outcomes (Rothman 1976). For any preventive or curative medical intervention, we can divide the sufficient causes of a health outcome Y into four types:

A: $M_T = 1,\ C_D = 0$
B: $M_T = 1,\ X = 0$
C: $M_{NT} = 1$
D: $C_H = 1,\ X = 1$

A, B, C, and D are types of sufficient cause, while $M_T$, $C_D$, $M_{NT}$, $C_H$, and X are types of component cause. Each sufficient cause is present for an individual when and only when all of its listed conditions are met. Each sufficient cause and each component cause can be represented by a dichotomous variable (0,1), where 1 represents individual-level "exposure" and 0 represents individual-level "non-exposure." X is a preventive or curative medical intervention, while Y is the health outcome it cures or prevents. $M_T$ and $M_{NT}$ are mechanisms or pathology that cause Y. $M_T$ represents mechanisms targeted by X, while $M_{NT}$ represents mechanisms not targeted by X. $C_D$ represents "disruption cofactors" that help X to disrupt $M_T$ and cure/prevent Y. $C_H$ represents "harm cofactors" that help X cause Y, a paradoxical harm. Exposure to $M_T$ or $M_{NT}$ occurs given at least one complete targeted mechanism or non-targeted mechanism, respectively. Exposure to $C_D$ or $C_H$ occurs given at least one complete set of disruption cofactors or harm cofactors, respectively.

We can combine the SCC model with the "potential outcomes" or "counterfactual framework" (Greenland and Robins 1986) to yield a combined counterfactual SCC model (Suzuki et al. 2018). In the counterfactual SCC model for a comparison involving exposure and non-exposure to a single intervention X, each individual in the population has potential outcomes for Y under exposure to X ($Y_{X=1}$) and under non-exposure to X ($Y_{X=0}$), which are determined by the sufficient causes present and absent for that individual. (In the absence of all of its sufficient causes, the outcome does not occur.) Assuming determinism (Greenland and Robins 1986) and given sufficient causes A–D, the potential outcome equations are:

$Y_{X=0} = \max(A,\ B,\ C)$

$Y_{X=1} = \max(A,\ C,\ D)$

The absolute risk in the population under non-exposure to X is $p(Y_{X=0})$.
The absolute risk under exposure to X is $p(Y_{X=1})$.
The absolute risk reduction (ARR) is given by $ARR = p(Y_{X=0}) - p(Y_{X=1})$.
The relative risk reduction (RRR) is given by $RRR = \left[p(Y_{X=0}) - p(Y_{X=1})\right] / p(Y_{X=0})$.

We can specify these equations for different scenarios in which only some of the four types of sufficient cause are present in the population, as in example scenarios 1–3 below. Scenario 4 modifies scenario 1 by relaxing the assumption of determinism and introducing two types of targeted mechanisms.

## 1. No Non-Targeted Mechanisms, No Paradoxical Harms

In this scenario, only targeted mechanisms ($M_T$) are present in the population; non-targeted mechanisms ($M_{NT}$) are not. Moreover, X never causes Y in the population; it only ever prevents Y. An example would be a non-live virus vaccine against a specific infection (e.g., X is a flu vaccine, Y is influenza). Then only sufficient causes A and B are present.

A: $M_T = 1,\ C_D = 0$

B: $M_T = 1,\ X = 0$

The absolute risks, ARR, and RRR can be calculated as follows:

(1) $p(Y_{X=0}) = p(M_T = 1)$

(2) $p(Y_{X=0}) = p(M_T = 1,\ C_D = 0) + p(M_T = 1,\ C_D = 1)$

(3) $p(Y_{X=1}) = p(M_T = 1,\ C_D = 0)$

$$ARR = p(Y_{X=0}) - p(Y_{X=1})$$

From (2) and (3):

$$ARR = p(M_T = 1,\ C_D = 1)$$

(4) $$\mathbf{ARR = p(M_T = 1) \bullet p(C_D = 1 |\ M_T = 1)}$$

$$RRR = \left[p(Y_{X=0}) - p(Y_{X=1})\right] / p(Y_{X=0})$$

From (1) and (4):

$$RRR = \frac{p(M_T = 1) \bullet p(C_D = 1 |\ M_T = 1)}{p(M_T = 1)}$$

(5) $$\mathbf{RRR = p(C_D = 1 | M_T = 1)}$$

In this scenario, the RRR is determined by (i) the probability of disruption cofactors among those with a targeted mechanism. The ARR is determined by the former probability, as well as (ii) the probability of a targeted mechanism.

## 2. Non-Targeted Mechanisms, No Paradoxical Harms

In this scenario, targeted mechanisms ($M_T$) are present in the population, as are non-targeted mechanisms ($M_{NT}$). Moreover, the intervention never causes Y in the population; it only ever prevents Y. An example would be a non-live virus vaccine against a non-specific infection outcome (e.g., X is a flu vaccine, Y is flu-like illness). Then only sufficient causes A, B, and C are present.

A: $M_T = 1,\ C_D = 0$

B: $M_T = 1,\ X = 0$

C: $M_{NT} = 1$

The absolute risks, ARR, and RRR can be calculated as follows:

(1) $p(Y_{X=0}) = p(M_T = 1,\ M_{NT} = 0) + p(M_{NT} = 1)$

(2) $p(Y_{X=0}) = p(M_T = 1,\ C_D = 0,\ M_{NT} = 0) + p(M_T = 1,\ C_D = 1,\ M_{NT} = 0) + p(M_{NT} = 1)$

(3) $p(Y_{X=1}) = p(M_T = 1,\ C_D = 0,\ M_{NT} = 0) + p(M_{NT} = 1)$

$$ARR = p(Y_{X=0}) - p(Y_{X=1})$$

From (2) and (3):

$$ARR = p(M_T = 1,\ C_D = 1,\ M_{NT} = 0)$$

(4) $$\mathbf{ARR = p(M_T = 1,\ M_{NT} = 0) \bullet p(C_D = 1 |\ M_T = 1,\ M_{NT} = 0)}$$

$$RRR = \left[p(Y_{X=0}) - p(Y_{X=1})\right] / p(Y_{X=0})$$

From (1) and (4):

(5) $$\mathbf{RRR = \frac{p(M_T = 1, M_{NT} = 0) \bullet p(C_D = 1 | M_T = 1, M_{NT} = 0)}{p(M_T = 1,\ M_{NT} = 0) + p(M_{NT} = 1)}}$$

Similar to the first scenario, the ARR is determined by (i) the probability of disruption cofactors among those with a targeted mechanism (and without a non-targeted mechanism) and (ii) the probability of a targeted mechanism (in the absence of a non-targeted mechanism). Unlike in the first scenario, the RRR is determined by the former two probabilities, as well as (iii) the probability of a non-targeted mechanism.

## 3. No Non-Targeted Mechanisms, Paradoxical Harms

In this scenario, only targeted mechanisms ($M_T$) are present in the population; non-targeted mechanisms ($M_{NT}$) are not. Moreover, X sometimes causes Y in the population, and other times prevents Y. A (rare) example would be a live virus vaccine against a specific infection (e.g., X is an attenuated poliovirus vaccine, Y is poliovirus infection). Then only sufficient causes A, B and D are present.

A: $M_T = 1, C_D = 0$

B: $M_T = 1,\ X = 0$

D: $C_H = 1,\ X = 1$

The absolute risks, ARR, and RRR can be calculated as follows:

(1) $p(Y_{X=0}) = p(M_T = 1)$

(2) $p(Y_{X=0}) = p(M_T = 1,\ C_D = 0) + p(M_T = 1,\ C_D = 1)$

(3) $p(Y_{X=1}) = p(M_T = 1,\ C_D = 0) + p(C_H = 1, A = 0)$

$$ARR = p(Y_{X=0}) - p(Y_{X=1})$$

From (2) and (3):

$$ARR = p(M_T = 1,\ C_D = 1) - p(C_H = 1, A = 0)$$

(4) $$\mathbf{ARR = p(M_T = 1) \cdot p(C_D = 1 |\ M_T = 1) - p(C_H = 1, A = 0)}$$

$$RRR = \left[p(Y_{X=0}) - p(Y_{X=1})\right] / p(Y_{X=0})$$

From (1) and (4):

$$RRR = \frac{p(M_T = 1) \cdot p(C_D = 1 |\ M_T = 1) - p(C_H = 1, A = 0)}{p(M_T = 1)}$$

(5) $$\mathbf{RRR = p(C_D = 1 | M_T = 1) - \frac{p(C_H = 1, A = 0)}{p(M_T = 1)}}$$

Unlike in the previous two scenarios, both the RRR and the ARR are determined by the same three probabilities: (i) the probability of disruption cofactors among those with a targeted mechanism, (ii) the probability of a targeted mechanism, and (iii) the probability of harm cofactors (in the absence of an undisrupted targeted mechanism).

## 4. Scenario 1 with Two Types of Chancy Targeted Mechanisms

This scenario modifies scenario 1 by relaxing the assumption of determinism and introducing two types of targeted mechanisms, $M_{Ta}$ and $M_{Tb}$. The probability of Y in an individual given an undisrupted $M_{Ta}$ is *a*, while the probability of Y in an individual given an undisrupted $M_{Tb}$ is *b*. (The probability of Y given a disrupted mechanism remains zero.)

We can modify equations (1), (3), (4), and (5) from scenario 1 to calculate the absolute risks, ARR, and RRR as follows:

(1) $p(Y_{X=0}) = a \cdot p(M_{Ta} = 1) + b \cdot p(M_{Tb} = 1)$

(3) $p(Y_{X=1}) = a \cdot p(M_{Ta} = 1, C_D = 0) + b \cdot p(M_{Tb} = 1, C_D = 0)$

(4) $$\mathbf{ARR} = a \cdot \mathbf{p(M_{Ta} = 1) \cdot p(C_D = 1 | M_{Ta} = 1)} + b \cdot \mathbf{p(M_{Tb} = 1) \cdot p(C_D = 1 | M_{Tb} = 1)}$$

If $a = b$:

$$\mathbf{ARR} = a \cdot \mathbf{p(M_T = 1) \cdot p(C_D = 1 | M_T = 1)}$$

$$\text{(5)} \qquad \mathbf{RRR} = \frac{a \bullet \mathbf{p}(\mathbf{M}_{\mathbf{Ta}} = \mathbf{1}) \bullet \mathbf{p}(\mathbf{C}_{\mathbf{D}} = \mathbf{1} \,|\, \mathbf{M}_{\mathbf{Ta}} = \mathbf{1}) + b \bullet \mathbf{p}(\mathbf{M}_{\mathbf{Tb}} = \mathbf{1}) \bullet \mathbf{p}(\mathbf{C}_{\mathbf{D}} = \mathbf{1} \,|\, \mathbf{M}_{\mathbf{Tb}} = \mathbf{1})}{a \bullet \mathbf{p}(\mathbf{M}_{\mathbf{Ta}} = \mathbf{1}) + b \bullet \mathbf{p}(\mathbf{M}_{\mathbf{Tb}} = \mathbf{1})}$$

If $a = b$ or if $p(C_D = 1 | M_{Ta} = 1) = p(C_D = 1 | M_{Tb} = 1)$:

$$\mathbf{RRR} = \mathbf{p}(\mathbf{C}_{\mathbf{D}} = \mathbf{1} \,|\, \mathbf{M}_{\mathbf{T}} = \mathbf{1})$$

In this scenario, if *a* is equal to *b*, then the RRR is determined by (i) the probability of disruption cofactors among those with a targeted mechanism. The ARR is determined by the former probability, as well as (ii) the probability of a targeted mechanism and (iii) the value of *a*.

If $p(C_D = 1 | M_{Ta} = 1)$ is equal to $p(C_D = 1 | M_{Tb} = 1)$, then the RRR is again determined by (i) the probability of disruption cofactors among those with a targeted mechanism. The ARR is determined by (i)–(iv) below.

Otherwise, both the RRR and the ARR are determined by (i) the probability of each type of targeted mechanism, (ii) the probability of disruption cofactors among those with each type of targeted mechanism, (iii) the value of *a*, and (iv) the value of *b*.

# References


3M. 2022. *State of Science Index 2022: Global Report.* The 3M Company. https://multimedia.3m.com/mws/media/2183175O/3m-state-of-science-index-sosi-2022-global-report.pdf.

Abdool Karim, S. S., and N. Devnarain. 2022. "Time to Stop Using Ineffective Covid-19 Drugs." *New England Journal of Medicine* 387 (7): 654–655.

Abrams, Marshall. 2017. "Probability and Chance in Mechanisms." In *The Routledge Handbook of Mechanisms and Mechanical Philosophy*, edited by Stuart Glennan and Phyllis Illari. Routledge.

Achinstein, Peter. 2001. *The Book of Evidence*. Oxford University Press.

Ackerknecht, Erwin Heinz. 1967. *Medicine at the Paris Hospital, 1794–1848*. Johns Hopkins University Press.

Ackerknecht, Erwin H. 2016. *A Short History of Medicine*. Johns Hopkins University Press.

Adam, D. 2020. "Special Report: The Simulations Driving the World's Response to COVID-19." *Nature* 580 (7803): 316–318.

Adebamowo, C., O. Bah-Sow, F. Binka, et al. 2014. "Randomised Controlled Trials for Ebola: Practical and Ethical Issues." *Lancet* 384 (9952): 1423–1424.

Alberg, Anthony J., Malcolm V. Brock, Jean G. Ford, Jonathan M. Samet, and Simon D. Spivack. 2013. "Epidemiology of Lung Cancer: Diagnosis and Management of Lung Cancer, 3rd ed: American College of Chest Physicians Evidence-Based Clinical Practice Guidelines." *Chest* 143 (5 Suppl): e1S–e29S. https://doi.org/10.1378/chest.12-2345.

American Diabetes Association. 2021a. "Cardiovascular Disease and Risk Management: Standards of Medical Care in Diabetes—2021." *Diabetes Care* 44 (Suppl. 1): S125–S150.

American Diabetes Association. 2021b. "Introduction: Standards of Medical Care in Diabetes—2021." *Diabetes Care* 44 (Suppl. 1): S1–S2.

American Diabetes Association. 2021c. "Pharmacologic Approaches to Glycemic Treatment: Standards of Medical Care in Diabetes—2021." *Diabetes Care* 44 (Suppl. 1): S111–S124.

American Diabetes Association. 2021d. "1. Improving Care and Promoting Health in Populations: Standards of Medical Care in Diabetes—2021." *Diabetes Care* 44 (Suppl. 1): S7–S14.

American Heart Association. 1973. *Coronary Risk Handbook: Estimating Risk of Coronary Heart Disease in Daily Practice*. American Heart Association.

American Psychiatric Association. 2013. *Diagnostic and Statistical Manual of Mental Disorders: DSM-5*. American Psychiatric Publishing.

American Psychiatric Association. 2022. *Diagnostic and Statistical Manual of Mental Disorders: DSM-5TR*. American Psychiatric Publishing.

American Thoracic Society. 2000. "Diagnostic Standards and Classification of Tuberculosis in Adults and Children." *American Journal of Respiratory and Critical Care Medicine* 161 (4): 1376–1395. https://doi.org/10.1164/ajrccm.161.4.16141.

Andersen, Holly. 2012. "Mechanisms: What Are They Evidence for in Evidence-Based Medicine?" *Journal of Evaluation in Clinical Practice* 18 (5): 992–999. https://doi.org/10.1111/j.1365-2753.2012.01906.x.

Anderson, T. J., J. Gregoire, G. J. Pearson, et al. 2016. "2016 Canadian Cardiovascular Society Guidelines for the Management of Dyslipidemia for the Prevention of Cardiovascular Disease in the Adult." *Canadian Journal of Cardiology* 32 (11): 1263–1282.

Andreoletti, M. 2018. "More Than One Way to Measure? A Casuistic Approach to Cancer Clinical Trials." *Perspectives in Biology and Medicine* 61 (2): 174–190.

Andrews, N., J. Stowe, F. Kirsebom, et al. 2022. "Covid-19 Vaccine Effectiveness Against the Omicron (B.1.1.529) Variant." *New England Journal of Medicine* 386 (16): 1532–1546.

Angell, Marcia. 2005. *The Truth About the Drug Companies: How They Deceive Us and What to Do About It*. Random House.

Angus, Derek C., and Chung-Chou H. Chang. 2021. "Heterogeneity of Treatment Effect: Estimating How the Effects of Interventions Vary Across Individuals." *JAMA* 326 (22): 2312–2313. https://doi.org/10.1001/jama.2021.20552.

Anjum, Rani Lill, Samantha Copeland, and Elena Rocca. 2020a. "Medical Scientists and Philosophers Worldwide Appeal to EBM to Expand the Notion of 'Evidence.'" *British Medical Journal Evidence-Based Medicine* 25 (1): 6.

Anjum, Rani Lill, Samantha Copeland, and Elena Rocca. 2020b. *Rethinking Causality, Complexity and Evidence for the Unique Patient*. Springer.

Antman, Elliott M., and Joseph Loscalzo. 2015. 'ST-Segment Elevation Myocardial Infarction." In *Harrison's Principles of Internal Medicine*, edited by Dennis L. Kasper, Stephen L. Hauser, J. Larry Jameson, Anthony S. Fauci, Dan L. Longo, and Joseph Loscalzo. McGraw-Hill Education.

Armitage, P., and R. Doll. 1954. "The Age Distribution of Cancer and a Multi-Stage Theory of Carcinogenesis." *British Journal of Cancer* 8 (1): 1–12.

Armitage, P., and Richard Doll. 1961. "Stochastic Models for Carcinogenesis." In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, edited by Jerzy Neyman. University of California Press.

Armstrong, D. M., C. B. Martin, and U. T. Place. 1996. *Dispositions: A Debate*. Routledge.

Arnaud, Sarah, and Anne-Marie Gagné-Julien. 2024. "The New Self-advocacy Activism in Psychiatry: Toward a Scientific Turn." *Philosophical Psychology* 37 (8): 2623–2646. https://doi.org/10.1080/09515089.2023.2174425.

Arnett, D. K., R. S. Blumenthal, M. A. Albert, et al. 2019. "2019 ACC/AHA Guideline on the Primary Prevention of Cardiovascular Disease: A Report of the American College of Cardiology/American Heart Association Task Force on Clinical Practice Guidelines." *Circulation* 140 (11): e596–e646.

Aronowitz, Robert. 1998. *Making Sense of Illness: Science, Society and Disease*. Cambridge University Press.

Aronowitz, Robert A. 2007. *Unnatural History: Breast Cancer and American Society*. Cambridge University Press.

Aronowitz, Robert. 2015. *Risky Medicine: Our Quest to Cure Fear and Uncertainty*. University of Chicago Press.

Ashley, David J. B. 1969. "Colonic Cancer Arising in Polyposis Coli." *Journal of Medical Genetics* 6 (4): 376–378.

Baden, L. R., H. M. El Sahly, B. Essink, et al. 2021. "Efficacy and Safety of the mRNA-1273 SARS-CoV-2 Vaccine." *New England Journal of Medicine* 384 (5): 403–416. https://doi.org/10.1056/NEJMoa2035389.

Baetu, Tudor M. 2016. "The 'Big Picture': The Problem of Extrapolation in Basic Research." *British Journal for the Philosophy of Science* 67 (4): 941–964. https://doi.org/10.1093/bjps/axv018.

Baigent, C., A. Keech, P. M. Kearney, et al. 2005. "Efficacy and Safety of Cholesterol-Lowering Treatment: Prospective Meta-Analysis of Data from 90,056 Participants in 14 Randomised Trials of Statins." *Lancet* 366 (9493): 1267–1278. https://doi.org/10.1016/s0140-6736(05)67394-1.

Balshem, Howard, Mark Helfand, Holger J. Schünemann, et al. 2011. "GRADE Guidelines: 3. Rating the Quality of Evidence." *Journal of Clinical Epidemiology* 64 (4): 401–406.

Bareinboim, Elias, and Judea Pearl. 2013. "A General Algorithm for Deciding Transportability of Experimental Results." *Journal of Causal Inference* 1 (1): 107–134.

Barillot, Emmanuel, Laurence Calzone, Philippe Hupe, Jean-Philippe Vert, and Andrei Zinovyev. 2020. *Computational Systems Biology of Cancer*. CRC Press.

Barnes, P. J. 2012. "254. Asthma." In *Harrison's Principles of Internal Medicine*, edited by D. L. Longo. McGraw-Hill.

Barnett, H. J. M., D. W. Taylor, R. B. Haynes, et al. 1991. "Beneficial Effect of Carotid Endarterectomy in Symptomatic Patients with High-Grade Carotid Stenosis." *New England Journal of Medicine* 325 (7): 445–453.

Barnett, M. L., M. Gaye, A. B. Jena, and A. Mehrotra. 2022. "Association of County-Level Prescriptions for Hydroxychloroquine and Ivermectin with County-Level Political Voting Patterns in the 2020 US Presidential Election." *Journal of the American Medical Association Internal Medicine* 182 (4): 452–454.

Bartholow, Roberts. 1876. "The Degree of Certainty in Therapeutics." *Transactions of the Medical and Chirurgical Faculty of the State of Maryland*: 33–50.

Bartholow, Roberts. 1881. *On the Antagonism Between Medicines and Between Remedies and Diseases*. D. Appleton.

Bateman, E. D., S. S. Hurd, P. J. Barnes, et al. 2008. "Global Strategy for Asthma Management and Prevention: GINA Executive Summary." *European Respiratory Journal* 31 (1): 143–78. https://doi.org/10.1183/09031936.00138707.

Beaver, J. A., L. J. Howie, L. Pelosof, et al. 2018. "A 25-Year Experience of US Food and Drug Administration Accelerated Approval of Malignant Hematology and Oncology Drugs and Biologics: A Review." *Journal of the American Medical Association Oncology* 4: 849–856.

Bechtel, William. 2018. "The Importance of Constraints and Control in Biological Mechanisms: Insights from Cancer Research." *Philosophy of Science* 85 (4): 573–593. https://doi.org/10.1086/699192.

Bechtel, William, and Robert C. Richardson. 2010. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. MIT Press.

Beebee, Helen, and David Papineau. 1997. "Probability as a Guide to Life." *Journal of Philosophy* 94 (5): 217.

Beets, Michael W., R. Glenn Weaver, John P. A. Ioannidis, et al. 2020. "Identification and Evaluation of Risk of Generalizability Biases in Pilot Versus Efficacy/Effectiveness Trials: A Systematic Review and Meta-analysis." *International Journal of Behavioral Nutrition and Physical Activity* 17 (1): 19. https://doi.org/10.1186/s12966-020-0918-y.

Bernard, Claude. 1865. *An Introduction to the Study of Experimental Medicine*. Macmillan.

Bernard, Claude. 1875. "De l'emploi des moyennes en physiologie experimentale a propos de l'influence de l'effeuillage des betteraves sur la production." *Comptes rendus de l'Academie des Sciences* 81: 702–703.

Bernard, Claude. 1947. *Principes de Medicine Experimentale*. Presses Universitaires de France.

Bertalanffy, L. von. 1968. *General System Theory*. Braziller.

Bertolaso, Marta. 2016. *Philosophy of Cancer: A Dynamic and Relational View*. Springer.

Bertolaso, Marta, and John Dupré. 2018. "A Processual Perspective on Cancer." In *Everything Flows: Towards a Processual Philosophy of Biology*, edited by Daniel J. Nicholson and John Dupré. Oxford University Press.

Bhimraj, A., R. L. Morgan, A. H. Shumaker, et al. 2023. "Infectious Diseases Society of America Guidelines on the Treatment and Management of Patients with COVID-19." Version 11.0.0. Infectious Diseases Society of America. https://www.idsociety.org/practice-guideline/covid-19-guideline-treatment-and-management/.

Binney, N. 2018. "The Function of the Heart Is Not Obvious." *Studies in History and Philosophy of Biological and Biomedical Sciences* 68–69: 56–69. https://doi.org/10.1016/j.shpsc.2018.05.003.

Bird, Alexander. 1998. "Dispositions and Antidotes." *Philosophical Quarterly* 48: 227–234.

Bird, Alexander. 2019. "Systematicity, Knowledge, and Bias: How Systematicity Made Clinical Medicine a Science." *Synthese* 196 (3): 863–879.

Bishop, J. M. 1995. "Cancer: The Rise of the Genetic Paradigm." *Genes & Development* 9 (11): 1309–1315. https://doi.org/10.1101/gad.9.11.1309.

Blasimme, A., P. Maugeri, and P. L. Germain. 2013. "What Mechanisms Can't Do: Explanatory Frameworks and the Function of the p53 Gene in Molecular Oncology." *Studies in History and Philosophy of Biological and Biomedical Sciences* 44 (3): 374–384. https://doi.org/10.1016/j.shpsc.2013.02.001.

Bliss, Michael. 2011. *The Making of Modern Medicine: Turning Points in the Treatment of Disease*. University of Toronto Press.

Bliss, Michael. 2013. *The Discovery of Insulin*. University of Chicago Press.

Bluhm, R. 2005. "From Hierarchy to Network—A Richer View of Evidence for Evidence-Based Medicine." *Perspectives in Biology and Medicine* 48 (4): 535–547. https://doi.org/10.1353/pbm.2005.0082.

Bluhm, R., and K. Borgerson. 2018. "An Epistemic Argument for Research-Practice Integration in Medicine." *Journal of Medicine and Philosophy* 43 (4): 469–484.

Bluhm, Robyn, and Kirstin Borgerson. 2011. "Evidence-Based Medicine." In *Handbook of the Philosophy of Science*. Vol. 16: *Philosophy of Medicine*, edited by D. M. Gabbay, P. Thagard, and J. Woods. Elsevier.

Bolie, V. W. 1960. "Coefficients of Normal Blood Glucose Regulation." *Journal of Applied Physiology* 16: 783–788.

Bonner, Carissa, Lyndal J. Trevena, Wolfgang Gaissmaier, et al. 2021. "Current Best Practice for Presenting Probabilities in Patient Decision Aids: Fundamental Principles." *Medical Decision Making* 41 (7): 821–833.

Boorse, Christopher. 1977. "Health as a Theoretical Concept." *Philosophy of Science* 44 (4): 542–573.

Boorse, Christopher. 2011. "Concepts of Health and Disease." In *Handbook of the Philosophy of Science*. Vol. 16: *Philosophy of Medicine*, edited by D. M. Gabbay, P. Thagard, and J. Woods. Elsevier.

Borén, J., M. J. Chapman, R. M. Krauss, et al. 2020. "Low-Density Lipoproteins Cause Atherosclerotic Cardiovascular Disease: Pathophysiological, Genetic, and Therapeutic Insights: A Consensus Statement from the European Atherosclerosis Society Consensus Panel." *European Heart Journal* 41 (24): 2313–2330. https://doi.org/10.1093/eurheartj/ehz962.

Borgerson, Kirstin. 2009. "Valuing Evidence Bias and the Evidence Hierarchy of Evidence-Based Medicine." *Perspectives in Biology and Medicine* 52 (2): 218–233.

Bose, S., H. Clevers, and X. Shen. 2021. "Promises and Challenges of Organoid-Guided Precision Medicine." *Med* 2 (9): 1011–1026.

Bothwell, L. E., J. A. Greene, S. H. Podolsky, and D. S. Jones. 2016. "Assessing the Gold Standard—Lessons from the History of RCTs." *New England Journal of Medicine* 374 (22): 2175–2181.

Boutron, I., M. J. Page, J. P. T. Higgins, D. G. Altman, A. Lundh, and A. Hróbjartsson. 2024. "Considering Bias and Conflicts of Interest Among the Included Studies." In *Cochrane Handbook for Systematic Reviews of Interventions*, version 6.5, edited by J. P. T. Higgins, J. Thomas, J. Chandler, et al. Cochrane. www.training.cochrane.org/handbook.

Boyd, Richard. 1991. "Realism, Anti-Foundationalism and the Enthusiasm for Natural Kinds." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 61 (1–2): 127–148. http://www.jstor.org/stable/4320174.

Bradburn, Norman N., Nancy Cartwright, and Jonathan Fuller. 2017. "A Theory of Measurement." In *Measurement in Medicine: Philosophical Essays on Assessment and Evaluation*, edited by Leah McClimans. Rowman & Littlefield.

Braillard, P. A. 2010. "Systems Biology and the Mechanistic Framework." *History & Philosophy of the Life Sciences* 32 (1): 43–62.

Broadbent, Alex. 2009. "Causation and Models of Disease in Epidemiology." *Studies in History and Philosophy of Biological and Biomedical Sciences* 40 (4): 302–311.

Broadbent, Alex. 2011. "Inferring Causation in Epidemiology: Mechanisms, Black Boxes, and Contrasts." In *Causality in the Sciences*, edited by Phyllis McKay Illari, Federica Russo, and Jon Williamson. Oxford Scholarship Online.

Broadbent, Alex. 2013. *Philosophy of Epidemiology*. Palgrave Macmillan.

Broadbent, Alex. 2014. "Disease as a Theoretical Concept: The Case of 'HPV-itis.'" *Studies in History and Philosophy of Biological and Biomedical Sciences* 48 (Pt. B): 250–257. https://doi.org/10.1016/j.shpsc.2014.07.010.

Broadbent, Alex. 2019. *Philosophy of Medicine*. Oxford University Press.

Broadbent, Alex, ed. 2025. *The Oxford Handbook of Philosophy of Medicine*. Oxford University Press.

Broadbent, Alex, and Jonathan Fuller. 2020. "Philosophy of Medicine: A Dedicated Journal for an Emerging Field." *Philosophy of Medicine* 1 (1): 1–5.

Bunker, J. P. 1995. "Medicine Matters After All." *Journal of the Royal College of Physicians of London* 29 (2): 105–112.

Bunker, J. P., H. S. Frazier, and F. Mosteller. 1994. "Improving Health: Measuring Effects of Medical Care." *Milbank Quarterly* 72 (2): 225–258.

Burnham, John C. 2012. "The Death of the Sick Role." *Social History of Medicine* 25 (4): 761–776.

Burnham, John C. 2015. *Health Care in America: A History*. Johns Hopkins University Press.

Butler, D. 2008. "Translational Research: Crossing the Valley of Death." *Nature* 453 (7197): 840–842.

Bynum, W. F. 1994. *Science and the Practice of Medicine in the Nineteenth Century*. Cambridge University Press.

Bynum, W. F. 2006. "The Rise of Science in Medicine, 1850–1913." In *The Western Medical Tradition: 1800 to 2000*, edited by W. F. Bynum, Anne Hardy, Stephen Jacyna,Christopher Lawrence and E. M. (Tilli) Tansey. Cambridge University Press.

Canadian Task Force on the Periodic Health Examination. 1979. "The Periodic Health Examination: Canadian Task Force on the Periodic Health Examination." *Canadian Medical Association Journal* 121 (9): 1193–1254. https://pubmed.ncbi.nlm.nih.gov/115569.

Candido dos Reis, Francisco J., Gordon C. Wishart, Ed M. Dicks, et al. 2017. "An Updated PREDICT Breast Cancer Prognostication and Treatment Benefit Prediction Model with Independent Validation." *Breast Cancer Research* 19 (1): 58.

Canguilhem, Georges. 1966. *Le Normal et le Pathologique*. Presses Universitaires de France.

Caplan, A. L., C. Plunkett, and B. Levin. 2015. "Selecting the Right Tool for the Job." *American Journal of Bioethics* 15 (4): 4–10.

Carel, Havi. 2016. *Phenomenology of Illness*. Oxford University Press.

Carnap, Rudolf. 1945. "On Inductive Logic." *Philosophy of Science* 12 (2): 72–97.

Carpenter, Daniel. 2010. *Reputation and Power: Organizational Image and Pharmaceutical Regulation at the FDA*. Princeton University Press.

Carter, Codell. 2003. *The Rise of Causal Concepts of Disease*. Ashgate.

Cartwright, Nancy. 2007. "Are RCTs the Gold Standard?" *BioSocieties* 2 (1): 11–20. https://doi.org/10.1017/s1745855207005029.

Cartwright, Nancy. 2009. "Evidence-Based Policy: What's to Be Done About Relevance?" *Philosophical Studies* 143 (1): 127–136. https://doi.org/10.1007/s11098-008-9311-4. https://doi.org/10.1007/s11098-008-9311-4.

Cartwright, Nancy. 2010. "What are Randomised Controlled Trials Good for?" *Philosophical Studies* 147 (1): 59–70. https://doi.org/10.1007/s11098-009-9450-2.

Cartwright, Nancy. 2011. "A Philosopher's View of the Long Road from RCTs to Effectiveness." *Lancet* 377 (9775): 1400–1401. https://doi.org/10.1016/s0140-6736(11)60563-1.

Cartwright, Nancy. 2021. "Rigour Versus the Need for Evidential Diversity." *Synthese* 199 (5): 13095–13119.

Cartwright, Nancy, and Jeremy Hardie. 2012. *Evidence-Based Policy: A Practical Guide to Doing It Better*. Oxford University Press.

Case, Anne, and Angus Deaton. 2020. *Deaths of Despair and the Future of Capitalism*. Princeton University Press.

Catalogue of Bias Collaboration, B. Holman, L. Bero, and B Mintzes. 2019. "Industry Sponsorship Bias." Catalogue of Bias. https://catalogofbias.org/biases/industry-sponsorship-bias/.

Chalmers, Iain. 2005. "Statistical Theory Was Not the Reason That Randomization Was Used in the British Medical Research Council's Clinical Trial of Streptomycin for Pulmonary Tuberculosis." In *Body Counts: Medical Quantification in Historical and Sociological Perspectives*, edited by Gerard Jorland, Annick Opinel, and George Weisz. McGill–Queen's University Press.

Chapin, Charles V. 1885. *The Present State of the Germ Theory of Disease*. Kellogg Printing Co.

Chapman, Robert. 2019. "Neurodiversity Theory and Its Discontents: Autism, Schizophrenia, and the Social Model of Disability." In *The Bloomsbury Companion to Philosophy of Psychiatry*, edited by Serife Tekin and Robyn Bluhm. Bloomsbury Academic.

Chin-Yee, Benjamin. 2024. "Cancer Medicine and Precision Oncology." In *Handbook of the Philosophy of Medicine*, edited by Thomas Schramme and Mary Walker. Springer.

Chou, R., A. Cantor, T. Dana, et al. 2022. "Statin Use for the Primary Prevention of Cardiovascular Disease in Adults: Updated Evidence Report and Systematic Review for the US Preventive Services Task Force." *Journal of the American Medical Association* 328 (8): 754–771.

Christensen, D. 2010. "Higher-Order Evidence." *Philosophy and Phenomenological Research* 81 (1): 185–215.

Cipriani, A., T. A. Furukawa, G. Salanti, et al. 2018. "Comparative Efficacy and Acceptability of 21 Antidepressant Drugs for the Acute Treatment of Adults with Major Depressive Disorder: A Systematic Review and Network Meta-Analysis." *Lancet* 391 (10128): 1357–1366.

Clarke, Brendan, Donald Gillies, Phyllis Illari, Federica Russo, and Jon Williamson. 2014. "Mechanisms and the Evidence Hierarchy." *Topoi* 33: 339.

Clarke, T. C., L. I. Black, B. J. Stussman, P. M. Barnes, and R. L. Nahin. 2015. "Trends in the Use of Complementary Health Approaches Among Adults: United States, 2002-2012." *National Health Statistics Report* 79: 1–16.

Cochrane, Archibald. 1972. *Effectiveness and Efficiency: Random Reflections on Health Services*. Nuffield Trust.

Cogliano, V. J., R. A. Baan, K. Straif, Y. Grosse, B. Secretan, and F. E. Ghissassi. 2008. "Use of Mechanistic Data in IARC Evaluations." *Environmental and Molecular Mutagenesis* 49: 100–109.

Colgrove, J. 2002. "The McKeown Thesis: A Historical Controversy and Its Enduring Influence." *American Journal of Public Health* 92 (5): 725–729.

Collins, Francis S. 2021. "Latest on Omicron Variant and COVID-19 Vaccine Protection." *NIH Director's Blog* (blog), December 14. https://directorsblog.nih.gov/2021/12/14/the-latest-on-the-omicron-variant-and-vaccine-protection/.

Collins, Francis S., and Harold Varmus. 2015. "A New Initiative on Precision Medicine." *New England Journal of Medicine* 372 (9): 793–795.

Collins, R., L. Bowman, M. Landray, and R. Peto. 2020. "The Magic of Randomization Versus the Myth of Real-World Evidence." *New England Journal of Medicine* 382 (7): 674–678.

Colnet, B., J. Josse, G. Varoquaux, et al. 2024. "Risk Ratio, Odds Ratio, Risk Difference . . . Which Causal Measure Is Easier to Generalize?" March 30. arXiv:2303.16008v3.

Comte, Auguste. 2012. *Cours de philosophie positive, leçons 46–51*. Hermann.

Conley, Brandon A. 2023. "Dispositionalism and Dysfunction." *Philosophy of Science* 90 (3): 686–703.

Connolly, K. R., and M. E. Thase. 2011. "If at First You Don't Succeed: A Review of the Evidence for Antidepressant Augmentation, Combination and Switching Strategies." *Drugs* 71 (1): 43–64. https://doi.org/10.2165/11587620-000000000-00000.

Conrad, Lawrence I., Michael Neve, Vivian Nutton, Roy Porter, and Andrew Wear, eds. 1995. *The Western Medical Tradition: 800 BC to AD 1800*. Cambridge University Press.

Cooper, Rachel. 2014. *Psychiatry and Philosophy of Science*. Routledge.

Cooter, Roger. 2011. "Medicine and Modernity." In *The Oxford Handbook of the History of Medicine*, edited by Mark Jackson. Oxford University Press.

Copi, Irving M., Carl Cohen, and Kenneth McMahon. 2014. *Introduction to Logic*. 14th ed. Pearson Education Limited.

Cornfield, Jerome. 1951. "A Method of Estimating Comparative Rates from Clinical Data: Applications to Cancer of the Lung, Breast, and Cervix." *Journal of the National Cancer Institute* 11 (6): 1269–1275.

Cornfield, Jerome, William Haenszel, E. Cuyler Hammond, Abraham M. Lilienfeld, Michael B. Shimkin, and Ernst L. Wynder. 1959. "Smoking and Lung Cancer: Recent Evidence and a Discussion of Some Questions." *Journal of the National Cancer Institute* 22 (1): 173–203. https://doi.org/10.1093/jnci/22.1.173.

Costa, J., M. Borges, C. David, et al. 2006. "Efficacy of Lipid Lowering Drug Treatment for Diabetic and Non-diabetic Patients: Meta-analysis of Randomised Controlled Trials." *BMJ* 332 (7550): 1115–1124. https://doi.org/10.1136/bmj.38793.468449.AE.

Coussens, Lisa M., and Zena Werb. 2002. "Inflammation and Cancer." *Nature* 420 (6917): 860–867. https://doi.org/10.1038/nature01322.

Cox, Daniel A., M. Anthony Mills, Ian R. Banks, Kelsey Eyre Hammond, and Kyle Gray. 2023. *America's Crisis of Confidence: Rising Mistrust,Conspiracies, and Vaccine Hesitancy After COVID-19*. Report. American Enterprise Institute, Washington, DC. https://www.americansurveycenter.org/research/americas-crisis-of-confidence-rising-mistrust-conspiracies-and-vaccine-hesitancy-after-covid-19/.

Cuello-Garcia, Carlos A., Nancy Santesso, Rebecca L. Morgan, et al. 2022. "GRADE Guidance 24 Optimizing the Integration of Randomized and Non-Randomized Studies of Interventions in Evidence Syntheses and Health Guidelines." *Journal of Clinical Epidemiology* 142: 200–208.

Curtis, H. J., P. Inglesby, C. E. Morton, et al. 2022. "Trends and Clinical Characteristics of COVID-19 Vaccine Recipients: A Federated Analysis of 57.9 Million Patients' Primary Care Records In Situ Using OpenSAFELY." *British Journal of General Practitioners* 72 (714): e51–62.

Daly, Jeanne. 2005. *Evidence-Based Medicine and the Search for a Science of Clinical Care*. University of California Press.

Dammann, Olaf. 2020. *Etiological Explanations*. CRC Press.

Dans, A. L., L. F. Dans, G. H. Guyatt, and S. Richardson. 1998. "Users' Guides to the Medical Literature. XIV: How to Decide on the Applicability of Clinical Trial Results to Your Patient." *Journal of the American Medical Association* 279 (7): 545–549.

Dawber, R., W. B. Kannel, and P. McNamara. 1964. "The Prediction of Coronary Heart Disease." In *Transactions of the Association of Life Insurance Medical Directors of America, 72nd Annual Meeting*, edited by S. R. Moore. Recording & Statistical Company.

Dawber, T. R., G. F. Meadors, and F. E. Moore Jr. 1951. "Epidemiological Approaches to Heart Disease: The Framingham Study." *American Journal of Public Health Nations Health* 41 (3): 279–281.

Dawid, Philip. 2017. "On Individual Risk." *Synthese* 194 (9): 3445–3474.

Dean, Michael Emmans. 2001. "The Trials of Homeopathy: A Critical-Historical Account of the Origins, Structure and Development of Hahnemann's Scientific Therapeutics, and Two Systematic Reviews of Homeopathic Clinical Trials, 1821–1953 and 1940–1998." PhD diss., University of York.

Deaton, Angus, and Nancy Cartwright. 2018. "Understanding and Misunderstanding Randomized Controlled Trials." *Social Science & Medicine* 210: 2–21.

Deeks, J. J. 2002. "Issues in the Selection of a Summary Statistic for Meta-Analysis of Clinical Trials with Binary Outcomes." *Statistics in Medicine* 21 (11): 1575–1600.

Dembić, Sanja. 2023. "Mental Disorder: An Ability-Based View." *Philosophy and the Mind Sciences* 4 (2): 1–28. https://doi.org/10.33735/phimisci.2023.9630.

de Melo-Martín, I. 2019. "The Commercialization of the Biomedical Sciences: (Mis)understanding Bias." *History & Philosophy of the Life Sciences* 41 (3): 34.

Denny, J. C., J. L. Rutter, D. B. Goldstein, et al. 2019. "The 'All of Us' Research Program." *New England Journal of Medicine* 381 (7): 668–676.

Desai, Angel. 2020. "Twentieth-Century Lessons for a Modern Coronavirus Pandemic." *Journal of the American Medical Association* 323 (21): 2118–2119.

DesAutels, L. 2015. "Toward a Propensity Interpretation of Stochastic Mechanism for the Life Sciences." *Synthese* 192 (9): 2921–2953. doi:10.1007/s11229-015-0694-4.

Devanesan, Arjun. 2020. "Medical Nihilism: The Limits of a Decontextualised Critique of Medicine." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 79: 101189.

Devanesan, Arjun. 2023. "The Reference Class Problem and Probabilities in the Individual Case: A Response to Fuller." *Philosophy of Science* 90 (4): 1001–1009.

Dickersin, K., and F. Chalmers. 2015. "Thomas C. Chalmers (1917–1995): A Pioneer of Randomised Clinical Trials and Systematic Reviews." *Journal of the Royal Society of Medicine* 108 (6): 237–241.

Dimaras, Helen, Timothy W. Corson, David Cobrinik, et al. 2015. "Retinoblastoma." *Nature Reviews Disease Primers* 1 (1): 15021. https://doi.org/10.1038/nrdp.2015.21.

Djulbegovic, B., C. L. Bennett, and G. H. Lyman. 1999. "Violation of the Uncertainty Principle in Conduct of Randomized Controlled Trials (RCTs) of Erythropoietin (EPO)." *Blood* 94 (Suppl. 1): 399A.

Djulbegovic, B., and G. H. Guyatt. 2017. "Progress in Evidence-Based Medicine: A Quarter Century On." *Lancet* 390 (10092): 415–423. https://doi.org/10.1016/s0140-6736(16)31592-6.

Djulbegovic, Benjamin, Gordon H. Guyatt, and Richard E. Ashcroft. 2009. "Epistemologic Inquiries in Evidence-Based Medicine." *Cancer Control* 16 (2): 158–168.

Djulbegovic, B., and I. Hozo. 2007. "When Should Potentially False Research Findings Be Considered Acceptable?" *PLoS Medicine* 4 (2): e26.

Djulbegovic, Benjamin, Iztok Hozo, and Sander Greenland. 2011. "Uncertainty in Clinical Medicine." In *Handbook of the Philosophy of Science*. Vol. 16: *Philosophy of Medicine*, edited by D. M. Gabbay, P. Thagard, and J. Woods. Elsevier.

Djulbegovic, B., A. Kumar, B. Miladinovic, et al. 2013. "Treatment Success in Cancer: Industry Compared to Publicly Sponsored Randomized Controlled Trials." *PLoS One* 8 (3): e58711. https://doi.org/10.1371/journal.pone.0058711.

Djulbegovic, B., M. Lacevic, A. Cantor, et al. 2000. "The Uncertainty Principle and Industry-Sponsored Research." *Lancet* 356 (9230): 635–638.

Doll, Richard. 1992. "Sir Austin Bradford Hill and the Progress of Medical Science." *BMJ* 305 (6868): 1521–1526. https://doi.org/10.1136/bmj.305.6868.1521.

Doll, Richard. 1994. "Austin Bradford Hill, 8 July 1897–18 April 1991." *Biographical Memoirs of Fellows of the Royal Society* 40: 127–140. https://doi.org/https://doi.org/10.1098/rsbm.1994.0032.

Doll, Richard. 2002. "Proof of Causality: Deduction from Epidemiological Observation." Perspectives in Biology and Medicine 45 (4): 499–515. https://doi.org/10.1353/pbm.2002.0067.

Doll, Richard, and Austin Bradford Hill. 1950. "Smoking and Carcinoma of the Lung: Preliminary Report." *British Medical Journal* 2 (4682): 739–748. https://doi.org/10.1136/bmj.2.4682.739.

Doll, Richard, and Austin Bradford Hill. 1954. "The Mortality of Doctors in Relation to Their Smoking Habits: A Preliminary Report." *British Medical Journal* 4877: 1451–1455.

Doll, Richard, and Austin Bradford Hill. 1956. "Lung Cancer and Other Causes of Death in Relation to Smoking: A Second Report on the Mortality of British Doctors." *British Medical Journal* 2 (5001): 1071–1081.

Dominus, Susan. 2020. "The Covid Drug Wars That Pitted Doctor vs. Doctor." *New York Times*, August 8.

Dorresteijn, J. A., F. L. Visseren, P. M. Ridker, et al. 2011. "Estimating Treatment Effects for Individual Patients Based on the Results of Randomised Clinical Trials." *British Medical Journal* 343: d5888.

Douglas, H. 2000. "Inductive Risk and Values in Science." *Philosophy of Science* 67 (4): 559–579. https://doi.org/10.1086/392855.

Drolet, B. C., and N. M. Lorenzi. 2011. "Translational Research: Understanding the Continuum from Bench to Bedside." *Translational Research* 157 (1): 1–5.

Dubos, R. 1959. *Mirage of Health: Utopias, Progress, and Biological Change*. Rutgers University Press.

Duffin, Jacalyn. 2005. *Lovers and Livers: Disease Concepts in History*. University of Toronto Press.

Early Breast Cancer Trialists' Collaborative Group. 2021. "Trastuzumab for Early-Stage, HER2-Positive Breast Cancer: A Meta-Analysis of 13,864 Women in Seven Randomised Trials." *Lancet Oncology* 22 (8): 1139–1150.

Ederer, Fred. 1982. "Jerome Cornfield's Contributions to the Conduct of Clinical Trials." *Biometrics* 38: 25–32.

Eersel, Gerdien G. van, Gabriela V. Koppenol-Gonzalez, and Julian Reiss. 2019. "Extrapolation of Experimental Results Through Analogical Reasoning from Latent Classes." *Philosophy of Science* 86 (2): 219–235. https://doi.org/10.1086/701956.

Ehrlich, Paul. 1908. "Experimental Researches on Specific Therapy: On Immunity with Special Reference to the Relationship Between Distribution and Action of Antigens: First Harben Lecture." In *The Harben Lectures for 1907 of the Royal Institute of Public Health*, edited by Paul Ehrlich. Lewis.

Elliott, Kevin C. 2014. "Financial Conflicts of Interest and Criteria for Research Credibility." *Erkenntnis* 79 (5): 917–937.

Emmady, P. D., C. Schoo, and P. Tadi. 2022. "Major Neurocognitive Disorder (Dementia)." StatPearls. https://www.ncbi.nlm.nih.gov/books/NBK557444/.

Engel, George L. 1977. "The Need for a New Medical Model—A Challenge for Biomedicine." *Science* 196 (4286): 129–136.

Engel, G. L. 1981. "The Clinical Application of the Biopsychosocial Model." *Journal of Medicine and Philosophy* 6 (2): 101–23.

Engelhard, Kristina. 2021. "Dispositions Manifest Themselves: An Identity Theory of Properties." *Synthese* 199 (5): 13497–13522.

Erasmus, Adrian. 2023. "The Bias Dynamics Model: Correcting for Meta-biases in Therapeutic Prediction." *Philosophy of Science* 90 (5): 1204–1213. https://doi.org/10.1017/psa.2023.66.

Ernst, Edzard. 2016. *Homeopathy: The Undiluted Facts.* Springer.

European Carotid Surgery Trialists' Collaborative Group. 1991. "MRC-European-Carotid-Surgery-Trial—Interim Results for Symptomatic Patients with Severe (70–99-Percent) or with Mild (0–29-Percent) Carotid Stenosis." *Lancet* 337 (8752): 1235–1243.

Evans, Alfred S. 1976. "Causation and Disease: The Henle-Koch Postulates Revisited." *The Yale Journal of Biology and Medicine* 49 (2): 175–195.

Evans, Alfred S. 1993. *Causation and Disease: A Chronological Journey*. Plenum.

Evans, J. P., C. Zeng, C. Carlin, et al. 2022. "Neutralizing Antibody Responses Elicited by SARS-CoV-2 mRNA Vaccination Wane Over Time and Are Boosted by Breakthrough Infection." *Science Translational Medicine* 14 (637): eabn8057.

Every-Palmer, S., and J. Howick. 2014. "How Evidence-Based Medicine Is Failing Due to Biased Trials and Selective Publication." *Journal of Evaluation in Clinical Practice* 20 (6): 908–914.

Evidence-Based Medicine Working Group. 1992. "Evidence-Based Medicine: A New Approach to Teaching the Practice of Medicine." *Journal of the American Medical Association* 268 (17): 2420–2425.

Eyding, Dirk, Monika Lelgemann, Ulrich Grouven, et al. 2010. "Reboxetine for Acute Treatment of Major Depression: Systematic Review and Meta-Analysis of Published and Unpublished Placebo and Selective Serotonin Reuptake Inhibitor Controlled Trials." *British Medical Journal* 341: c4737.

Fagan, Melinda Bonnie. 2017. "Explanation, Multiple Perspectives, and Understanding." *Balkan Journal of Philosophy* 1: 19–34.

Fanu, James Le. 2019. *Too Many Pills: How Too Much Medicine Is Endangering Our Health and What We Can Do About It*. Abacus.

Fauci, A. S. 2022. "It Ain't Over Till It's Over . . . but It's Never Over—Emerging and Reemerging Infectious Diseases." *New England Journal of Medicine* 387 (22): 2009–2011.

Feldman, R. 2005. "Respecting the Evidence." *Philosophical Perspectives* 19: 95–119.

Feudtner, Chris. 2003. *Bittersweet: Diabetes, Insulin and the Transformation of Illness*. University of North Carolina Press.

Fisher, Ronald A. 1935. *The Design of Experiments*. Oliver & Boyd.

Fisher, R. A. 1958a. "Cancer and Smoking." *Nature* 182: 596.

Fisher, R. A. 1958b. "Cigarettes, Cancer, and Statistics." *Centennial Review* 2: 151–166.

Fleck, Ludwik. 1979. *Genesis and Development of a Scientific Fact*. University of Chicago Press.

Fletcher, Samuel C., Jürgen Landes, and Roland Poellinger. 2019. "Evidence Amalgamation in the Sciences: An Introduction." *Synthese* 196 (8): 3163–3188. https://doi.org/10.1007/s11229-018-1840-6. https://doi.org/10.1007/s11229-018-1840-6.

Flexner, Abraham. 1910. *Medical Education in the United States and Canada*. Carnegie Foundation for the Advancement of Teaching.

Flexner, Abraham. 1925. *Medical Education: A Comparative Study*. Macmillan.

Food and Drug Administration. 2013. *Paving the Way for Personalized Medicine: FDA's Role in a New Era of Medical Product Development*. Food and Drug Administration.

Food and Drug Administration. 2018. "The Drug Development Process." https://www.fda.gov/patients/learn-about-drug-and-device-approvals/drug-development-process.

Food and Drug Administration. 2022. "Homeopathic Drug Products: Guidance for FDA Staff and Industry." U.S. Department of Health and Human Services.

Food and Drug Administration. 2024. "FDA Approves First Treatment to Reduce Risk of Serious Heart Problems Specifically in Adults with Obesity or Overweight." Press release, March 8. https://www.fda.gov/news-events/press-announcements/fda-approves-first-treatment-reduce-risk-serious-heart-problems-specifically-adults-obesity-or.

Forlino, Antonella, Wayne A. Cabral, Aileen M. Barnes, and Joan C. Marini. 2011. "New Perspectives on Osteogenesis Imperfecta." *Nature Reviews Endocrinology* 7 (9): 540–557. https://doi.org/10.1038/nrendo.2011.81.

Foucault, Michel. 1970. *The Order of Things: An Archaeology of the Human Sciences*. Pantheon Books.

Foucault, Michel. 1973. *The Birth of the Clinic: An Archaeology of Medical Perception*. Routledge.

Frank, Steven A. 2007. *Dynamics of Cancer: Incidence, Inheritance, and Evolution*. Princeton: Princeton University Press.

Frass, M., R. P. Strassl, H. Friehs, M. Müllner, M. Kundi, and A. D. Kaye. 2012. "Use and Acceptance of Complementary and Alternative Medicine Among the General Population and Medical Personnel: A Systematic Review." *Ochsner Journal* 12 (1): 45–56.

Frieden, T. R. 2017. "Evidence for Health Decision Making—Beyond Randomized, Controlled Trials." *New England Journal of Medicine* 377 (5): 465–475.

Friedman, Daniel C., and Dunja Šešelja. 2023. "Scientific Disagreements, Fast Science and Higher-Order Evidence." *Philosophy of Science* 90: 937–57.

Friend, Toby. 2021. "Megarian Variable Actualism." *Synthese* 199 (3): 10521–10541.

Frost, Wade Hampton. 1939. "The Age Selection of Mortality from Tuberculosis in Successive Decades." *American Journal of Epidemiology* 30 (3): 91–96. https://doi.org/10.1093/oxfordjournals.aje.a118570.

Frost, Wade Hampton. 1941. "Epidemiology." In *Hampton Frost, MD: A Contribution to Epidemiologic Method*, edited by Kenneth F. Maxcy. Commonwealth Fund.

Fujiki, Hirota, Eisaburo Sueoka, and Masami Suganuma. 2013. "Tumor Promoters: From Chemicals to Inflammatory Proteins." *Journal of Cancer Research and Clinical Oncology* 139 (10): 1603–1614. https://doi.org/10.1007/s00432-013-1455-8.

Fujimura, Joan H. 1988. "The Molecular Biological Bandwagon in Cancer Research: Where Social Worlds Meet." *Social Problems* 35 (3): 261–283. https://doi.org/10.2307/800622.

Fuller, Jonathan. 2013a. "Rationality and the Generalization of Randomized Controlled Trial Evidence." *Journal of Evaluation in Clinical Practice* 19 (4): 644–647.

Fuller, Jonathan. 2013b. "Rhetoric and Argumentation: How Clinical Practice Guidelines Think." *Journal of Evaluation in Clinical Practice* 19 (3): 433–441.

Fuller, Jonathan. 2018a. "Meta-Research Evidence for Evaluating Therapies." *Philosophy of Science* 85 (5): 767–780.

Fuller, Jonathan. 2018b. "Universal Etiology, Multifactorial Diseases and the Constitutive Model of Disease Classification." *Studies in History and Philosophy of Biological and Biomedical Sciences* 67: 8–15. https://doi.org/10.1016/j.shpsc.2017.11.002.

Fuller, Jonathan. 2018c. "What Are Chronic Diseases?" *Synthese* 195 (7): 3197–3220. https://doi.org/10.1007/s11229-017-1368-1.

Fuller, Jonathan. 2019. "The Confounding Question of Confounding Causes in Randomized Trials." *British Journal for the Philosophy of Science* 70 (3): 901–926.

Fuller, Jonathan. 2020a. "From Pandemic Facts to Pandemic Policies." In *Thinking in a Pandemic: The Crisis of Science and Policy in the Age of COVID-19*, edited by Matt Lord. Boston Review & Verso Books. Original publication, https://www.bostonreview.net/articles/jonathan-fuller-pandemic-facts-pandemic-policies/.

Fuller, Jonathan. 2020b. "Medical Nihilism by Jacob Stegenga: What Is the Right Dose?" *Studies in History and Philosophy of Biological and Biomedical Sciences* 81: 101270.

Fuller, Jonathan. 2020c. "Models Versus Evidence." In *Thinking in a Pandemic: The Crisis of Science and Policy in the Age of COVID-19*, edited by Matt Lord. Boston Review & Verso Books. Original publication, https://www.bostonreview.net/articles/jonathan-fuller-models-v-evidence/.

Fuller, Jonathan. 2021a. "Demarcating and Judging Medicine: Review of Broadbent's 'Philosophy of Medicine.'" *Philosophy of Science* 88 (2): 370–376.

Fuller, Jonathan. 2021b. "The Myth and Fallacy of Simple Extrapolation in Medicine." *Synthese* 198 (4): 2919–2939.

Fuller, Jonathan. 2022a. "Epidemics from the Population Perspective." *Philosophy of Science* 89 (2): 232–251.

Fuller, Jonathan. 2022b. "Preventive and Curative Medical Interventions." *Synthese* 200 (2): 61.

Fuller, Jonathan. 2024. "Demarcating Scientific Medicine." *Studies in History and Philosophy of Science* 106: 177–185.

Fuller, Jonathan. 2025. "Modern Scientific Medicine's Demarcation Criterion." In *The Oxford Handbook of Philosophy of Medicine*, edited by Alex Broadbent. Oxford University Press.

Fuller, Jonathan, Benjamin Chin-Yee, and Ross E. G. Upshur. 2024. "The Argument Framework Is a Flexible Approach to Evidence in Healthcare." *Nature Medicine* 30 (8): 2113–2116. https://doi.org/10.1038/s41591-024-02930-x.

Furukawa, T. A., G. H. Guyatt, and L. E. Griffith. 2002. "Can We Individualize the 'Number Needed to Treat'? An Empirical Study of Summary Effect Measures in Meta-Analyses." *International Journal of Epidemiology* 31 (1): 72–76.

Futcher, Thomas B. 1907. "Diabetes Mellitus." In *Modern Medicine: Its Theory and Practice*, edited by William Osler and Thomas McCrae. Lea Brothers.

Fye, W. Bruce. 1987. *The Development of American Physiology*. Johns Hopkins University Press.

Gadsbøll, N., P. F. Hoilundcarlsen, G. G. Nielsen, et al. 1989. "Symptoms and Signs of Heart-Failure in Patients with Myocardial-Infarction—Reproducibility and Relationship to Chest-X-Ray, Radionuclide Ventriculography and Right Heart Catheterization." *European Heart Journal* 10 (11): 1017–1028.

Gale, E. A. 2002. "The Rise of Childhood Type 1 Diabetes in the 20th Century." *Diabetes* 51 (12): 3353–3361.

Gamma, A. 2013. "The Role of Genetic Information in Personalized Medicine." *Perspectives in Biology and Medicine* 56 (4): 485–512.

Gamma, Alex. 2017. "Personalized and Precision Medicine." In *The Routledge Companion to Philosophy of Medicine*, edited by Miriam Solomon, Jeremy Simon, and Harold Kinkaid. Routledge.

Gaudillière, Jean-Paul. 2002. *Inventer la biomédecine: La France, l'Amérique et la production des savoirs du vivant (1945–1965)*. La Découverte.

Giere, Ronald N. 1999. *Science Without Laws*. University of Chicago Press.

Gifford, Fred. 2011. "Philosophy of Medicine: Introduction." In *Handbook of the Philosophy of Science*. Vol. 16: *Philosophy of Medicine*, edited by D. M. Gabbay, P. Thagard, and J. Woods. Elsevier.

Gifford, Fred. 2017. "The Biomedical Model and the Biopsychosocial Model." In *The Routledge Companion to Philosophy of Medicine*, edited by Miriam Solomon, Jeremy R. Simon, and Harold Kincaid. Routledge.

Gigerenzer, G., and A. Edwards. 2003. "Simple Tools for Understanding Risks: From Innumeracy to Insight." *British Medical Journal* 327 (7417): 741–744. https://doi.org/10.1136/bmj.327.7417.741.

Gillies, Donald. 2000. *Philosophical Theories of Probability*. Routledge.

Gillies, Donald. 2005. "Hempelian and Kuhnian Approaches in the Philosophy of Medicine: The Semmelweis Case." *Studies in History and Philosophy of Biological and Biomedical Sciences* 36 (1): 159–81.

Gillies, Donald. 2019. "Should We Distrust Medical Interventions?" *Metascience* 28 (2): 273–276.

Giroux, Élodie. 2011. "The Origins of the Prospective Cohort Study: American Cardiovascular Epidemiology and the Framingham Heart Study." *Revue d'histoire des sciences* 64 (2): 297–318.

Giroux, Élodie. 2012. "The Framingham Study and the Constitution of a Restrictive Concept of Risk Factor." *Social History of Medicine* 26 (1): 94–112. https://doi.org/10.1093/shm/hks051.

Giroux, Élodie. 2015. "Risk Factor and Causality in Epidemiology." In *Classification, Disease and Evidence: New Essays in the Philosophy of Medicine*, edited by Philippe Huneman, Gérard Lambert, and Marc Silberstein. Springer.

Glasziou, P. P., and L. M. Irwig. 1995. "An Evidence Based Approach to Individualising Treatment." *British Medical Journal* 311 (7016): 1356–1359.

Glasziou, Paul, and David Mant. 2007. "Applying Results to Treatment Decisions in Primary Care." In *Treating Individuals: From Randomized Trials to Personalised Medicine*, edited by Peter M. Rothwell, 83–95. *The Lancet*.

Glasziou, P. P., S. Sanders, and T. Hoffmann. 2020. "Waste in Covid-19 Research." *British Medical Journal* 369: m1847.

Global Burden of Disease. 2020. "Global Burden of Disease Study 2019 Results." Institute for Health Metrics and Evaluation. https://vizhub.healthdata.org/gbd-results/.

Global Burden of Disease Mental Disorders Collaborators. 2022. "Global, Regional, and National Burden of 12 Mental Disorders in 204 Countries and Territories, 1990–2019: A Systematic Analysis for the Global Burden of Disease Study 2019." *Lancet Psychiatry* 9 (2): 137–150.

Goff, D. C., Jr., D. M. Lloyd-Jones, G. Bennett, et al. 2014. "2013 ACC/AHA Guideline on the Assessment of Cardiovascular Risk: A Report of the American College of Cardiology/American Heart Association Task Force on Practice Guidelines." *Circulation* 129: S49–73.

Goldacre, Ben. 2012. *Bad Pharma: How Drug Companies Mislead Doctors and Harm Patients*. Fourth Estate.

Goldenberg, Maya J. 2021. *Vaccine Hesitancy: Public Trust, Expertise, and the War on Science*. University of Pittsburgh Press.

Gomes, B., and E. A. Ashley. 2023. "Artificial Intelligence in Molecular Medicine." *New England Journal of Medicine* 388 (26): 2456–2465.

Goodman, Nelson. 1954. *Fact, Fiction, and Forecast*. Harvard University Press.

Goodman, Richard A., S. F. Posner, E. S. Huang, A. K. Parekh, and H. K. Koh. 2013. "Defining and Measuring Chronic Conditions: Imperatives for Research, Policy, Program, and Practice." *Preventing Chronic Disease* 10: 1–16. doi:10.5888/pcd10.120239.

Goodman, Steven. 1999. "Probability at the Bedside: The Knowing of Chances or the Chances of Knowing?" *Annals of Internal Medicine* 130 (7): 604–606.

Goodman, Steven, and Kay Dickersin. 2011. "Metabias: A Challenge for Comparative Effectiveness Research." *Annals of Internal Medicine* 155 (1): 61.

Goodman, Steven, and Sander Greenland. 2007. "Why Most Published Research Findings Are False: Problems in the Analysis." *PLoS Medicine* 4 (4): e168.

Goodwin, William. 2015. "Revolution and Progress in Medicine." *Theoretical Medicine and Bioethics* 36: 25–39.

Gordin, Michael D. 2021. *On the Fringe: Where Science Meets Pseudoscience*. Oxford University Press.

Gøtzsche, Peter C. 2013. *Deadly Medicines and Organized Crime: How Big Pharma Has Corrupted Healthcare*. Radcliffe.

Govindaraju, D. R., L. Adrienne Cupples, W. B. Kannel, et al. 2008. "Genetics of the Framingham Heart Study Population." *Advanced Genetics* 62: 33–65.

Gradle, H. 1883. *Bacteria, and the Germ-Theory of Disease*. Keener.

Gray, J. A. 1999. "Postmodern Medicine." *Lancet* 354 (9189): 1550–1553.

Green, S., A. Carusi, and K. Hoeyer. 2019. "Plastic Diagnostics: The Remaking of Disease and Evidence in Personalized Medicine." *Social Science & Medicine* 304: 1–9.

Green, Sara, Mie Seest Dam, and Mette Nordahl Svendsen. 2022. "Patient-Derived Organoids in Precision Oncology: Towards a Science of and for the Individual?" In *Personalized Medicine in the Making*, edited by C. Beneduce and M. Bertolaso. Springer.

Green, Sara, Maria Şerban, Raphael Scholl, Nicholas Jones, Ingo Brigandt, and William Bechtel. 2018. "Network Analyses in Systems Biology: New Strategies for Dealing with Biological Complexity." *Synthese* 195 (4): 1751–1777. https://doi.org/10.1007/s11229-016-1307-6. https://doi.org/10.1007/s11229-016-1307-6.

Green, Sara, and Henrik Vogt. 2016. "Personalizing Medicine: Disease Prevention in Silico and in Socio." *Humana.Mente Journal of Philosophical Studies* 30: 105–145.

Greene, Jeremy. 2007. *Prescribing by Numbers: Drugs and the Definition of Disease*. Johns Hopkins University Press.

Greene, Jeremy A., Flurin Condrau, and Elizabeth Siegel Watkins. 2016. "Medicine Made Modern by Medicines." In *Therapeutic Revolutions: Pharmaceuticals and Social Change in the Twentieth Century*, edited by Jeremy A. Greene, Flurin Condrau, and Elizabeth Siegel Watkins. University of Chicago Press.

Greene, Jeremy A., David S. Jones, and Scott H. Podolsky. 2012. "Theraputic Evolution and the Challenge of Rational Medicine." *New England Journal of Medicine* 367 (12): 1077–1082.

Greene, Jeremy A., and Joseph Loscalzo. 2017. "Putting the Patient Back Together—Social Medicine, Network Medicine, and the Limits of Reductionism." *New England Journal of Medicine* 377 (25): 2493–2499. https://doi.org/10.1056/NEJMms1706744.

Greenhalgh, Trisha. 2010. *How to Read a Paper: The Basics of Evidence-Based Medicine*. 4th ed. John Wiley & Sons.

Greenhalgh, Trisha. 2020a. "Will COVID-19 Be Evidence-Based Medicine's Nemesis?" *PLoS Medicine* 17 (6): e1003266.

Greenhalgh, Trisha. 2020b. "Will Evidence-Based Medicine Survive COVID-19?" In *Thinking in a Pandemic: The Crisis of Science and Policy in the Age of COVID-19*, edited by Matt Lord. Boston Review & Verso Books. Original publication, https://www.bostonreview.net/articles/trisha-greenhalgh-covid-19-and-evidence-based-medicine/.

Greenhalgh, Trisha, David Fisman, Danielle J. Cane, Matthew Oliver, and Chandini Raina Macintyre. 2022. "Adapt or Die: How the Pandemic Made the Shift from EBM to EBM+ More Urgent." *British Medical Journal Evidence-Based Medicine* 27 (5): 253–260.

Greenland, S., and B. Brumback. 2002. "An Overview of Relations Among Causal Modelling Methods." *International Journal of Epidemiology* 31 (5): 1030–1037. https://doi.org/10.1093/ije/31.5.1030.

Greenland, S., J. Pearl, and J. M. Robins. 1999. "Causal Diagrams for Epidemiologic Research." *Epidemiology* 10 (1): 37–48.

Greenland, Sander, and James M. Robins. 1986. "Identifiability, Exchangeability, and Epidemiological Confounding." *International Journal of Epidemiology* 15 (3): 413–419.

Grob, Gerald N. 2005. *The Deadly Truth: A History of Disease in America*. Harvard University Press.

Gross, F. 2011. "What Systems Biology Can Tell Us About Disease." *History & Philosophy of the Life Sciences* 33 (4): 477–496.

Grundy, S. M., N. J. Stone, A. L. Bailey, et al. 2019. "2018 AHA/ACC/AACVPR/AAPA/ABC/ACPM/ADA/AGS/APhA/ASPC/NLA/PCNA Guideline on the Management of Blood Cholesterol: Executive Summary: A Report of the American College of Cardiology/American Heart Association Task Force on Clinical Practice Guidelines." *Journal of the American College of Cardiology* 73 (24): 3168–3209.

Guyatt, Gordon, Andrew Oxman, and Regina Kunz. 2008. "GRADE: An Emerging Consensus on Rating Quality of Evidence and Strength of Recommendations." *British Medical Journal* 336: 924–928.

Guyatt, Gordon H., Andrew D. Oxman, Regina Kunz, et al. 2011. "GRADE Guidelines: 8. Rating the Quality of Evidence—Indirectness." *Journal of Clinical Epidemiology* 64 (12): 1303–1310.

Guyatt, G. H., A. D. Oxman, V. Montori, et al. 2011. "GRADE Guidelines: 5. Rating the Quality of Evidence—Publication Bias." *Journal of Clinical Epidemiology* 64 (12): 1277–1282.

Guyatt, Gordon, Drummond Rennie, Maureen O. Meade, and Deborah J. Cook. 2015. *Users' Gudes to the Medical Literature: A Manual for Evidence-Based Clinical Practice*. 3rd ed. McGraw-Hill Education.

Gyawali, B., A. S. Kesselheim, and J. S. Ross. 2023. "The Accelerated Approval Program for Cancer Drugs—Finding the Right Balance." *New England Journal of Medicine* 389 (11): 968–971.

Hacking, Ian. 1965. *Logic of Statistical Inference*. Cambridge University Press.

Hacking, Ian. 2001. *An Introduction to Probability and Inductive Logic*. Cambridge University Press.

Hahnemann, Samuel. 1810. *Organon of the Healing Art*. Arnoldsche Buchhandlung.

Hajat, C., and E. Stein. 2018. "The Global Burden of Multiple Chronic Conditions: A Narrative Review." *Preventive Medicine Reports* 12: 284–293.

Hájek, Alan. 2007. "The Reference Class Problem Is Your Problem Too." *Synthese* 156 (3): 563–585.

Halm, E. A., S. Tuhrim, J. J. Wang, C. Rockman, T. S. Riles, and M. R. Chassin. 2009. "Risk Factors for Perioperative Death and Stroke After Carotid Endarterectomy: Results of the New York Carotid Artery Surgery Study." *Stroke* 40 (1): 221–229.

Hamburg, M. A., and F. S. Collins. 2010. "The Path to Personalized Medicine." *New England Journal of Medicine* 363 (4): 301–304.

Hammond, E. Cuyler. 1955. "Cause and Effect." In *The Biologic Effects of Tobacco*, edited by Ernest L. Wynder. Little, Brown.

Han, Paul. 2021. *Uncertainty in Medicine: A Framework for Tolerance*. Oxford University Press.

Han, P. K., T. C. Lehman, H. Massett, S. J. Lee, W. M. Klein, and A. N. Freedman. 2009. "Conceptual Problems in Laypersons' Understanding of Individualized Cancer Risk: A Qualitative Study." *Health Expectations* 12 (1): 4–17.

Hanahan, Douglas. 2022. "Hallmarks of Cancer: New Dimensions." *Cancer Discovery* 12 (1): 31–46.

Hanahan, D., and R. A. Weinberg. 2000. "The Hallmarks of Cancer." *Cell* 100 (1): 57–70. https://doi.org/10.1016/s0092-8674(00)81683-9.

Hanahan, D., and R. A. Weinberg. 2011. "Hallmarks of Cancer: The Next Generation." *Cell* 144 (5): 646–674. https://doi.org/10.1016/j.cell.2011.02.013.

Hankinson, R. J. 2016. "Humours and Humoral Theory." In *The Routledge History of Disease*, edited by M. Jackson. Routledge.

Hansen, C., A. Lundh, K. Rasmussen, and A. Hróbjartsson. 2019. "Financial Conflicts of Interest in Systematic Reviews: Associations with Results, Conclusions, and Methodological Quality." *Cochrane Database of Systematic Reviews* 8: 1–51.

Hansson, G. K. 2005. "Mechanisms of Disease—Inflammation, Atherosclerosis, and Coronary Artery Disease." *New England Journal of Medicine* 352 (16): 1685–1695. https://doi.org/10.1056/NEJMra043430. <Go to ISI>://WOS:000228515200010.

Haydock, Paul M., and Martin R. Cowie. 2010. "Heart Failure: Classification and Pathophysiology." *Medicine* 38 (9): 467–472.

Hayes, D. F. 2019. "HER2 and Breast Cancer—A Phenomenal Success Story." *New England Journal of Medicine* 381 (13): 1284–1286. https://doi.org/10.1056/NEJMcibr1909386.

Healy, David. 2012. *Pharmageddon*. University of California Press.

Heinbokel, Timm. 2024. "The Pragmatist Roots of Scientific Medicine: Reassessing Abraham Flexner's Report on Medical Education." *Studies in History and Philosophy of Science* 106: 186–195.

Hempel, Carl G., and Paul Oppenheim. 1948. "Studies in the Logic of Explanation." *Philosophy of Science* 15 (2): 135–175.

Hennekens, C. H. 2015. "Overview of Primary Prevention of Coronary Heart Disease and Stroke." *UpToDate*. https://www.uptodate.com/contents/overview-of-primary-prevention-of-cardiovascular-disease-in-adults.

Herrera-Perez, D., A. Haslam, T. Crain, et al. 2019. "A Comprehensive Review of Randomized Clinical Trials in Three Medical Journals Reveals 396 Medical Reversals." *Elife* 8: e45183.

Herrington, William, Ben Lacey, Paul Sherliker, Jane Armitage, and Sarah Lewington. 2016. "Epidemiology of Atherosclerosis and the Potential to Reduce the Global Burden of Atherothrombotic Disease." *Circulation Research* 118 (4): 535–546. https://doi.org/doi:10.1161/CIRCRESAHA.115.307611.

Hesslow, Germund. 1988. "The Problem of Causal Selection." In *Contemporary Science and Natural Explanation: Commonsense Conceptions of Causality*, edited by D. J. Hilton. New York University Press.

Hey, Spencer Phillips. 2015. "What Theories Are Tested in Clinical Trials?" *Philosophy of Science* 82 (5): 1318–1329. https://doi.org/10.1086/683816.

Higgins, Julian P. T., Douglas G. Altman, and Jonathan A. C. Sterne. 2011. "Assessing Risk of Bias in Included Studies." In *Cochrane Handbook for Systematic Reviews of Interventions*, version 5.1.0, edited by J. P. T. Higgins and S. Green. Cochrane Collaboration. www.handbook.cochrane.org.

Hill, Austin Bradford. 1937. *Principles of Medical Statistics*. The Lancet Limited.
Hill, Austin Bradford. 1952. "The Clinical Trial." *New England Journal of Medicine* 247: 113–119.
Hill, Austin Bradford. 1953. "Observation and Experiment." *New England Journal of Medicine* 248 (24): 995–1001.
Hill, Austin Bradford. 1965. "The Environment and Disease: Association or Causation?" *Proceedings of the Royal Society of Medicine* 58 (5): 295–300.
Hill, Austin Bradford, and I. D. Hill. 1991. *Bradford Hill's Principles of Medical Statistics*. 12th ed. Edward Arnold.
Hoffman, Jerome R., Michael S. Wilkes, Frank C. Day, Douglas S. Bell, and Jerilyn K. Higa. 2006. "The Roulette Wheel: An Aid to Informed Decision Making." *PLOS Medicine* 3 (6): e137.
Hogg, Jabez. 1870. "The Organic Germ Theory of Disease." *Medical Times and Gazette* 659: 685.
Holman, Bennett, and Kevin C. Elliott. 2018. "The Promise and Perils of Industry-Funded Science." *Philosophy Compass* 13 (11): e12544.
Holmes, Oliver Wendell. 1842. *Homeopathy and its Kindred Delusions*. William D. Ticknor.
Hopewell, S., K. Loudon, M. J. Clarke, A. D. Oxman, and K. Dickersin. 2009. "Publication Bias in Clinical Trials Due to Statistical Significance or Direction of Trial Results." *Cochrane Database of Systematic Reviews* 1: Mr000006.
Horby, P., M. Mafham, L. Linsell, et al. 2020. "Effect of Hydroxychloroquine in Hospitalized Patients with Covid-19." *New England Journal of Medicine* 383 (21): 2030–2040. https://doi.org/10.1056/NEJMoa2022926.
Horby, P., W. S. Lim, J. R. Emberson, et al. 2021. "Dexamethasone in Hospitalized Patients with Covid-19." *New England Journal of Medicine* 384 (8): 693–704. https://doi.org/10.1056/NEJMoa2021436.
Horn, L., C. M. Lovly, and D. H. Johnson. 2015. "Neoplasms of the Lung." In *Harrison's Principles of Internal Medicine*, edited by Stephen L. Hauser, Dennis L. Kasper, J. Larry Jameson, Anthony S. Fauci, Dan L. Longo and Joseph Loscalzo. McGraw-Hill.
Horowitz, S. 2014. "Epistemic Akrasia." *Nous* 48 (4): 718–744.
Horsburgh, C. R. 2015. "Epidemiology of Tuberculosis." *UpToDate*. https://www.uptodate.com/contents/epidemiology-of-tuberculosis/print.
Horwitz, Allan V., and Jerome C. Wakefield. 2012. *The Loss of Sadness: How Psychiatry Transformed Normal Sorrow into Depressive Disorder*. Oxford University Press.
Hotez, Peter J. 2023. *The Deadly Rise of Anti-Science: A Scientist's Warning*. Johns Hopkins University Press.
Howick, Jeremy. 2011. *The Philosophy of Evidence-Based Medicine*. Wiley-Blackwell.
Howick, J. 2016. "Aulus Cornelius Celsus and 'Empirical' and 'Dogmatic' Medicine." *Journal of the Royal Society of Medicine* 109 (11): 426–430.
Howick, J., P. Glasziou, and J. K. Aronson. 2013. "Problems with Using Mechanisms to Solve the Problem of Extrapolation." *Theoretical Medicine and Bioethics* 34 (4): 275–291.
Howick, Jeremy, Despina Koletsi, John P. A. Ioannidis, et al. 2022. "Most Healthcare Interventions Tested in Cochrane Reviews Are Not Effective According to High Quality Evidence: A Systematic Review and Meta-Analysis." *Journal of Clinical Epidemiology* 148: 160–169.
Howson, Colin, and Peter Urbach. 2006. *Scientific Reasoning: the Bayesian Approach*. 3rd ed. Open Court.
Hoyningen-Huene, Paul. 2013. *Systematicity: The Nature of Science*. Oxford University Press.
Hruban, Ralph H., and Will Linder. 2022. *A Scientific Revolution: Ten Men and Women Who Reinvented American Medicine*. Pegasus Books. https://www.kff.org/coronavirus-covid-19/poll-finding/kff-covid-19-vaccine-monitor-september-2023/.
Huang, S. 2011. "Systems Biology of Stem Cells: Three Useful Perspectives to Help Overcome the Paradigm of Linear Pathways." *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 366 (1575): 2247–2259. https://doi.org/10.1098/rstb.2011.0008.

Huang, S., and L. Hood. 2019. "Personalized, Precision, and N-of-One Medicine: A Clarification of Terminology and Concepts." *Perspectives in Biology and Medicine* 62 (4): 617–639.

Hucklenbroich, P. 2014. "'Disease Entity' as the Key Theoretical Concept of Medicine." *Journal of Medicine and Philosophy* 39 (6): 609–633. https://doi.org/10.1093/jmp/jhu040.

Hudis, C. A. 2007. "Trastuzumab-echanism of Action and Use in Clinical Practice." *New England Journal of Medicine* 357 (1): 39–51. https://doi.org/10.1056/NEJMra043186.

Hultcrantz, M., D. Rind, E. A. Akl, et al. 2017. "The GRADE Working Group Clarifies the Construct of Certainty of Evidence." *Journal of Clinical Epidemiology* 87: 4–13. https://doi.org/10.1016/j.jclinepi.2017.05.006.

Hutter, C., and J. C. Zenklusen. 2018. "The Cancer Genome Atlas: Creating Lasting Value beyond Its Data." *Cell* 173 (2): 283–285. https://doi.org/10.1016/j.cell.2018.03.042.

Hyman, S. E. 2010. "The Diagnosis of Mental Disorders: The Problem of Reification." *Annual Review of Clinical Psychology* 6: 155–179.

Illari, Phyllis McKay, and Jon Williamson. 2012. "What is a mechanism? Thinking about mechanisms across the sciences." *European Journal for Philosophy of Science* 2 (1): 119-–135. https://doi.org/10.1007/s13194-011-0038-2.

Illich, Ivan. 1976. *Medical Nemesis: The Expropriation of Health*. Pantheon Books.

Insel, T. R. 2014. "The NIMH Research Domain Criteria (RDoC) Project: Precision Medicine for Psychiatry." *American Journal of Psychiatry* 171 (4): 395–397.

International Agency for Research on Cancer, Working Group on the Evaluation of Carcinogenic Risks to Humans. 2012. "Arsenic, Metals, Fibres and Dusts." IARC Monographs on the Evaluation of Carcinogenic Risks to Humans, No. 100C. https://www.ncbi.nlm.nih.gov/books/NBK304368/.

Intersalt Cooperative Research Group. 1988. "Intersalt: an international study of electrolyte excretion and blood pressure. Results for 24 hour urinary sodium and potassium excretion. Intersalt Cooperative Research Group." *British Medical Journal (Clinical research ed.)* 297 (6644): 319–328. https://doi.org/10.1136/bmj.297.6644.319.

Ioannidis, J. P. A. 2005. "Why Most Published Research Findings Are False." *PLoS Medicine* 2 (8): 696–701.

Ioannidis, John P. A. 2020. "The Totality of the Evidence." In *Thinking in a Pandemic: The Crisis of Science and Policy in the Age of COVID-19*, edited by Matt Lord. Boston Review & Verso Books. Original publication, https://www.bostonreview.net/articles/john-p-ioannidis-totality-evidence/.

Ioannidis, J. P. A., D. Fanelli, D. D. Dunne, and S. N. Goodman. 2015. "Meta-Research: Evaluation and Improvement of Research Methods and Practices." *PloS Biology* 13 (10): e1002264.

Jansen, Ludger. 2007. "Tendencies and Other Realizables in Medical Information Sciences." *Monist* 90 (4): 534–554.

Jäntgen, Ina. 2023. "How to Measure Effect Sizes for Rational Decision-Making." *Philosophy of Science* 90 (5): 1183–1193.

Jewson, N. D. 1976. "The Disappearance of the Sick-Man from Medical Cosmology, 1770–1870." *Sociology* 10 (2): 225–244.

Jolie, Angelina. 2013. "My Medical Choice." *New York Times*, May 14.

Jones, David S. 2013. *Broken Hearts: The Tangled History of Cardiac Care*. Johns Hopkins University Press.

Jones, David S., and G. M. Oppenheimer. 2017. "If the Framingham Heart Study Did Not Invent the Risk Factor, Who Did?" *Perspectives in Biology and Medicine* 60 (2): 131–150. https://doi.org/10.1353/pbm.2017.0024.

Jones, David S., and S. H. Podolsky. 2015. "The History and Fate of the Gold Standard." *Lancet* 385 (9977): 1502–1503. https://doi.org/10.1016/s0140-6736(15)60742-5.

Jones, David S., S. H. Podolsky, and J. A. Greene. 2012. "The Burden of Disease and the Changing Task of Medicine." *New England Journal of Medicine* 366 (25): 2333–2338.

Jones, K. E., N. G. Patel, M. A. Levy, et al. 2008. "Global Trends in Emerging Infectious Diseases." *Nature* 451 (7181): 990–993. https://doi.org/10.1038/nature06536.

Jung, Richard G., Pietro Di Santo, Cole Clifford, et al. 2021. "Methodological Quality of COVID-19 Clinical Research." *Nature Communications* 12 (1): 943.

Justiz Vaillant, A. A., and P. G. Gulick. 2022. "HIV and AIDS Syndrome." StatPearls. https://www.ncbi.nlm.nih.gov/books/NBK534860/.

Kannel, W. B., T. R. Dawber, A. Kagan, N. Revotskie, and J. Stokes 3rd. 1961. "Factors of Risk in the Development of Coronary Heart Disease—Six Year Follow-Up Experience: The Framingham Study." *Annals of Internal Medicine* 55: 33–50.

Karlawish, J. 2010. "Desktop Medicine." *Journal of the American Medical Association* 304 (18): 2061–2062.

Kasper, Dennis L., Anthony S. Fauci, Stephen L. Hauser, Dan L. Longo, J. Larry Jameson, and Joseph Loscalzo, eds. 2015. *Harrison's Principles of Internal Medicine*. McGraw-Hill Education.

Keating, Peter, and Alberto Cambrosio. 2003. *Biomedical Platforms: Realigning the Normal and the Pathological in Late-Twentieth Century Medicine*. MIT Press.

Keating, Peter, and Alberto Cambrosio. 2004. "Does Biomedicine Entail the Successful Reduction of Pathology to Biology?" *Perspectives in Biology and Medicine* 47 (3): 357–371.

Keating, P., and A. Cambrosio. 2007. "Cancer Clinical Trials: The Emergence and Development of a New Style of Practice." *Bulletin of the History of Medicine* 81 (1): 197–223.

Keating, Peter, and Alberto Cambrosio. 2011. *Cancer on Trial: Oncology as a New Style of Practice*. University of Chicago Press.

Keating, P., A. Cambrosio, and N. C. Nelson. 2016. "'Triple Negative Breast Cancer': Translational Research and the (Re)Assembling of Diseases in Post-Genomic Medicine." *Studies in History and Philosophy of Biological and Biomedical Sciences* 59: 20–34.

Kelly, Robert. 2023. "Towards a Dispositionalist (and Unifying) Account of Addiction." *Theoretical Medicine and Bioethics* 44 (1): 21–40.

Kelly, Thomas. 2023. *Bias: A Philosophical Study*. Oxford University Press.

Kemp, C. D., and J. V. Conte. 2012. "The Pathophysiology of Heart Failure." *Cardiovascular Pathology* 21 (5): 365–371. https://doi.org/10.1016/j.carpath.2011.11.007.

Kendler, K. S. 2017. "DSM Disorders and Their Criteria: How Should They Inter-Relate?" *Psychological Medicine* 47 (12): 2054–2060.

Kendler, Kenneth S., and Josef Parnas. 2013. *Philosophical Issues in Psychiatry II: Nosology*. Oxford University Press.

Kendler, Kenneth S., and Josef Parnas. 2017. *Philosophical Issues in Psychiatry IV: Classification of Psychiatric Illness*. Oxford University Press.

Kennedy-Martin, Tessa, Sarah Curtis, Douglas Faries, Susan Robinson, and Joseph Johnston. 2015. "A Literature Review on the Representativeness of Randomized Controlled Trial Samples and Implications for the External Validity of Trial Results." *Trials* 16 (1): 495.

Kent, David. M., and Rodney A. Hayward. 2007. "Limitations of Applying Summary Results of Clinical Trials to Individual Patients: The Need for Risk Stratification." *JAMA* 298 (10): 1209–1212. https://doi.org/10.1001/jama.298.10.1209.

Kent, David M., Jason Nelson, Issa J. Dahabreh, Peter M. Rothwell, Douglas G. Altman, and Rodney A. Hayward. 2016. "Risk and Treatment Effect Heterogeneity: Re-Analysis of Individual Participant Data from 32 Large Clinical Trials." *International Journal of Epidemiology* 45 (6): 2075–2088.

Kent, David M., E. Steyerberg, and D. van Klaveren. 2018. "Personalized Evidence Based Medicine: Predictive Approaches to Heterogeneous Treatment Effects." *British Medical Journal* 363: k4245. https://doi.org/10.1136/bmj.k4245.

Kermack, William Ogilvy, A. G. McKendrick, and Gilbert Thomas Walker. 1927. "A Contribution to the Mathematical Theory of Epidemics." *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 115 (772): 700–721.

KFF COVID-19 Vaccine Monitor. 2023. "Partisanship Remains Key Predictor of Views of COVID-19, Including Plans to Get Latest COVID-19 Vaccine." September. https://www.kff.org/health-information-trust/kff-covid-19-vaccine-monitor-september-2023/.

Khalidi, Muhammad Ali. 2013. *Natural Categories and Human Kinds: Classification in the Natural and Social Sciences*. Cambridge University Press.
Khalidi, Muhammad Ali. 2021. "Etiological Kinds." *Philosophy of Science* 88 (1): 1–21. https://doi.org/10.1086/710020.
Khan, S. S., J. Coresh, M. J. Pencina, et al. 2023. "Novel Prediction Equations for Absolute Risk Assessment of Total Cardiovascular Disease Incorporating Cardiovascular-Kidney-Metabolic Health: A Scientific Statement from the American Heart Association." *Circulation* 148 (24): 1982–2004.
Khosrowi, Donal. 2019. "Extrapolation of Causal Effects: Hopes, Assumptions, and the Extrapolator's Circle." *Journal of Economic Methodology* 26 (1): 45–58.
Khosrowi, Donal. 2023. "Extrapolating from Experiments, Confidently." *European Journal for Philosophy of Science* 13 (2): 18.
Kincaid, H. 2008. "Do We Need Theory to Study Disease? Lessons from Cancer Research and Their Implications for Mental Illness." *Perspectives in Biology and Medicine* 51 (3): 367–378. https://doi.org/10.1353/pbm.0.0019.
King, Lester S. 1954. "What Is Disease?" *Philosophy of Science* 21 (3): 193–203.
King, Lester S. 1982. *Medical Thinking: A Historical Preface*. Princeton University Press.
King, Lester S. 1991. *Transformations in American Medicine: From Benjamin Rush to William Osler*. Johns Hopkins University Press.
Knight, David. 2009. *The Making of Modern Science*. Wiley.
Knudson, A. G., Jr. 1971. "Mutation and Cancer: Statistical Study of Retinoblastoma." *Proceedings of the National Academy of Sciences of the United States of America* 68 (4): 820–823.
Koch, Robert. 1882. "Die Atiologie der Tuberkulose." *Berliner Klinischen Wochenschrift* 15: 221–230.
Kravitz, R. L., N. Duan, and J. Braslow. 2004. "Evidence-Based Medicine, Heterogeneity of Treatment Effects, and the Trouble with Averages." *Milbank Quarterly* 82 (4): 661–687.
Kreso, Antonija, and John E Dick. 2014. "Evolution of the Cancer Stem Cell Model." *Cell Stem Cell* 14 (3): 275–291. https://doi.org/10.1016/j.stem.2014.02.006.
Krieger, N. 1994. "Epidemiology and the Web of Causation: Has Anyone Seen the Spider?" *Social Science & Medicine* 39 (7): 887–903. https://doi.org/10.1016/0277-9536(94)90202-x.
Krieger, Nancy. 2011. *Epidemiology and the People's Health*. Oxford University Press.
Krimsky, Sheldon. 2003. *Science in the Private Interest: Has the Lure of Profits Corrupted Biomedical Research?* Rowman & Littlefield.
Kripke, Saul. 1972. "Naming and Necessity." In *Semantics of Natural Language*, edited by D. Davidson and G. Harman. Reidel.
Kuhn, Thomas S. 2012. *The Structure of Scientific Revolutions*. 4th ed. University of Chicago Press.
Kuipers, Ernst J., William M. Grady, David Lieberman, et al. 2015. "Colorectal Cancer." *Nature Reviews Disease Primers* 1 (1): 15065. https://doi.org/10.1038/nrdp.2015.65.
Kumar, Vinay, Abul K. Abbas, and Jon C. Aster, eds. 2015. *Robbins and Cotran Pathologic Basis of Disease*. 9th ed. Elsevier Saunders.
La Caze, Adam. 2009. "Evidence-Based Medicine Must Be . . ." *Journal of Medicine and Philosophy* 34: 509–527.
La Caze, Adam. 2011. "The Role of Basic Science in Evidence-Based Medicine." *Biology & Philosophy* 26 (1): 81–98. https://doi.org/10.1007/s10539-010-9231-5.
La Caze, Adam, and Barbara Osimani, eds. 2020. *Uncertainty in Pharmacology*: Springer.
LaFollette, Hugh, and Niall Shanks. 1995. "Two Models of Models in Biomedical Research." *Philosophical Quarterly* 45 (179): 141–160.
Lai, Yurong, Xiaoyan Chu, Li Di, et al. 2022. "Recent Advances in the Translation of Drug Metabolism and Pharmacokinetics Science for Drug Discovery and Development." *Acta Pharmaceutica Sinica B* 12 (6): 2751–2777.

Laing, R. D., and A. Esterson. 1970. *Sanity, Madness, and the Family: Families of Schizophrenics*. Pelican.

Lakatos, Imre. 1976. "Falsification and the Methodology of Scientific Research Programmes." In *Can Theories Be Refuted? Essays on the Duhem-Quine Thesis*, edited by Sandra G. Harding. Springer Netherlands.

*Lancet*. 2020. "The Truth Is Out There, Somewhere." 396 (10247): 291.

Landes, Jürgen. 2020. "Variety of Evidence and the Elimination of Hypotheses." *European Journal for Philosophy of Science* 10 (2): 12.

Landes, Jürgen, Barbara Osimani, and Roland Poellinger. 2018. "Epistemology of Causal Inference in Pharmacology: Towards a Framework for the Assessment of Harms." *European Journal for Philosophy of Science* 8 (1): 3–49. https://doi.org/10.1007/s13194-017-0169-1.

Lane-Claypon, Janet E. 1926. "A Further Report On Cancer of the Breast: With Special Reference to Its Associated Antecedent Conditions." *Reports on Public Health and Medical Subjects*. No. 32. Ministry of Health.

Lange, Marc. 2007. "The End of Diseases." *Philosophical Topics* 35 (1–2): 265–292.

Laplane, Lucie. 2014. "Identifying Some Theories in Developmental Biology: The Case of the Cancer Stem Cell Theory." In *Towards a Theory of Development*, edited by Alessandro Minelli and Thomas Pradeu. Oxford University Press.

Laplane, Lucie. 2016. *Cancer Stem Cells: Philosophy and Therapies*. Harvard University of Press.

Laplane, Lucie. 2018. "Cancer Stem Cells Modulate Patterns and Processes of Evolution in Cancers." *Biology & Philosophy* 33 (3): 18.

Larroulet Philippi, Cristian. 2022. "There Is Cause to Randomize." *Philosophy of Science* 89 (1): 152–170.

Larson, H. J., L. Z. Cooper, J. Eskola, S. L. Katz, and S. Ratzan. 2011. "Addressing the Vaccine Confidence Gap." *Lancet* 378 (9790): 526–535.

Lassaunière, Ria, Charlotta Polacek, Anders Frische, et al. 2022. "Neutralizing Antibodies Against the SARS-CoV-2 Omicron Variant (BA.1) 1 to 18 Weeks After the Second and Third Doses of the BNT162b2 mRNA Vaccine." *Journal of the American Medical Association Network Open* 5 (5): e2212073–e2212073.

Lau, Joseph, John P. A. Ioannidis, Norma Terrin, Christopher H. Schmid, and Ingram Olkin. 2006. "The Case of the Misleading Funnel Plot." *British Medical Journal* 333 (7568): 597–600.

Laudan, Larry. 1981. "A Confutation of Convergent Realism." *Philosophy of Science* 48 (1): 19–49.

Laudan, Larry. 1983. "The Demise of the Demarcation Problem." In *Physics, Philosophy and Psychoanalysis: Essays in Honour of Adolf Grünbaum,* edited by R. S. Cohen and L. Laudan, 111–127. Springer Netherlands.

Lawrence, Christopher, and George Weisz, eds. 1998. *Greater Than the Parts: Holism in Biomedicine, 1920–1950*. Oxford University Press.

Lazarus, Jeffrey V., Katarzyna Wyka, Trenton M. White, et al. 2023. "A Survey of COVID-19 Vaccine Acceptance Across 23 Countries in 2022." *Nature Medicine* 29 (2): 366–375.

Lean, Oliver M. 2020. "Binding Specificity and Causal Selection in Drug Design." *Philosophy of Science* 87 (1): 70–90. https://doi.org/10.1086/706093.

Lee, Keekok. 2012. *Philosophical Foundations of Modern Medicine*: Palgrave Macmillan.

Le Fanu, James. 2012. *The Rise and Fall of Modern Medicine*. Revised ed. Basic Books.

Lemoine, Maël. 2015. "The Naturalization of the Concept of Disease." In *Classification, Disease and Evidence: New Essays in the Philosophy of Medicine*, edited by Gérard Lambert, Marc Silberstein, and Philippe Huneman. Springer.

Lemoine, M. 2017a. "Animal Extrapolation in Preclinical Studies: An Analysis of the Tragic Case of TGN1412." *Studies in History and Philosophy of Biological and Biomedical Sciences* 61: 35–45. https://doi.org/10.1016/j.shpsc.2016.12.004.

Lemoine, Maël. 2017b. "Neither from Words, nor from Visions: Understanding P-Medicine from Innovative Treatments." *Lato Sensu: Revue de la Société de philosophie des sciences* 4 (2): 12–23.

Leonelli, Sabina. 2024. "Is Data Science Transforming Biomedical Research? Evidence, Expertise, and Experiments in COVID-19 Science." *Philosophy of Science*: 91 (5): 1338–1348.

Letai, Anthony. 2017. "Functional Precision Cancer Medicine—Moving Beyond Pure Genomics." *Nature Medicine* 23 (9): 1028–1035.

Levin, J. M., L. A. Bukowski, J. A. Minson, and J. M. Kahn. 2023. "The Political Polarization of COVID-19 Treatments Among Physicians and Laypeople in the United States." *Proceedings of the National Academy of Sciences of the United States of America* 120 (7): e2216179120.

Lewis, D. 1973. "Causation." *Journal of Philosophy* 70: 556–567.

Lewis, David. 1980. "A Subjectivist's Guide to Objective Chance." In *Studies in Inductive Logic and Probability*, vol 2, edited by Richard C. Jeffrey. University of California Press.

Lewis, David. 1997. "Finkish Dispositions." *Philosophical Quarterly* 47: 143–158.

Lexchin, J., L. A. Bero, B. Djulbegovic, and O. Clark. 2003. "Pharmaceutical Industry Sponsorship and Research Outcome and Quality: Systematic Review." *British Medical Journal* 326 (7400): 1167–1170.

Light, R. W. 2015. "Disorders of the Pleura." In *Harrison's Principles of Internal Medicine*, edited by Stephen L. Hauser, Dennis L. Kasper, J. Larry Jameson, Anthony S. Fauci, Dan L. Longo and Joseph Loscalzo. McGraw-Hill.

Lilienfeld, A. M. 1959. "On the Methodology of Investigations of Etiologic Factors in Chronic Diseases: Some Comments." *Journal of Chronic Disease* 10 (1): 41–6. https://doi.org/10.1016/0021-9681(59)90016-5.

Lilienfeld, A. M. 1982. "The Fielding H. Garrison Lecture: Ceteris Paribus: The Evolution of the Clinical Trial." *Bulletin of the History of Medicine* 56 (1): 1–18.

Lindsay, Robert, and Felicia Cosman. 2015. "Osteoporosis." In *Harrison's Principles of Internal Medicine*, edited by Dennis L. Kasper, Stephen L. Hauser, J. Larry Jameson, Anthony S. Fauci, Dan L. Longo, and Joseph Loscalzo. McGraw-Hill Education.

Lipsitch, Marc. 2020. "Good Science Is Good Science." In *Thinking in a Pandemic: The Crisis of Science and Policy in the Age of COVID-19*, edited by Matt Lord. Boston Review & Verso Books. Original publication, https://www.bostonreview.net/articles/marc-lipsitch-good-science-good-science/.

Lipton, Peter. 2004. *Inference to the Best Explanation*. 2nd ed. Routledge.

Ljungman, Per. 2012. "Vaccination of Immunocompromised Patients." *Clinical Microbiology and Infection* 18: 93–99.

London, A. J., and J. Kimmelman. 2020. "Against Pandemic Research Exceptionalism." *Science* 368 (6490): 476–477.

Longo, Dan L., ed. 2012. *Harrison's Principles of Internal Medicine*. 18th ed. McGraw-Hill.

Loudon, I. 2006. "A Brief History of Homeopathy." *Journal of the Royal Society of Medicine* 99 (12): 607–610.

Louis, Pierre Charles Alexandre. 1834. *An Essay on Clinical Instruction*. Translated by Peter Martin. Highley.

Louis, P. C. A. 1836. *Researches on the Effects of Bloodletting in Some Inflammatory Diseases: And on the Influence of Tartarized Antimony and Vesication in Pneumonitis*. Hillard, Gray.

Louis, Pierre Charles Alexandre. 1837. "De l'examen des malades et de la recherche des faits generaux." In *Memoires de la Societe Medicale d'Observation*, edited by P. C. A. Louis. Crochard.

Löwy, Ilana. 2011. "Historiography of Biomedicine: 'Bio,' 'Medicine,' and In Between." Isis 102 (1): 116–122.

Lundh, A., J. Lexchin, B. Mintzes, J. B. Schroll, and L. Bero. 2017. "Industry Sponsorship and Research Outcome." *Cochrane Database of Systematic Reviews* 2 (2): Mr000033.

Lüönd, Fabiana, Stefanie Tiede, and Gerhard Christofori. 2021. "Breast Cancer as an Example of Tumour Heterogeneity and Tumour Cell Plasticity During Malignant Progression." *British Journal of Cancer* 125 (2): 164–175. https://doi.org/10.1038/s41416-021-01328-7.

Lyon, Aidan. 2017. "Vague Credence." *Synthese* 194 (10): 3931–3954.

Machamer, P., L. Darden, and C. F. Craver. 2000. "Thinking About Mechanisms." *Philosophy of Science* 67 (1): 1–25. https://doi.org/10.1086/392759.

Mackenbach, Johan P. 2022. "Omran's 'Epidemiologic Transition' 50 Years On." *International Journal of Epidemiology* 51 (4): 1054–1057.

Mackie, J. L. 1965. "Causes and Conditions." *American Philosophical Quarterly* 2 (4): 245–264.

Mackie, J. L. 1980. *The Cement of the Universe: A Study of Causation*. Oxford Scholarship Online.

MacLagan, Thomas J. 1876a. "Correspondence." *Lancet* 107 (2738): 295–296.

MacLagan, Thomas J. 1876b. *The Germ Theory Applied to the Explanation of the Phenomena of Disease: The Specific Fevers*. Macmillan.

MacMahon, B., T. F. Pugh, and J. Ipsen. 1960. *Epidemiologic Methods*. Little, Brown.

Malaterre, Christophe. 2007. "Organicism and Reductionism in Cancer Research: Towards a Systemic Approach." *International Studies in the Philosophy of Science* 21 (1): 57–73. https://doi.org/10.1080/02698590701305792.

Mangione, C. M., M. J. Barry, W. K. Nicholson, et al. 2022. "Statin Use for the Primary Prevention of Cardiovascular Disease in Adults: US Preventive Services Task Force Recommendation Statement." *Journal of the American Medical Association* 328 (8): 746–753.

Mann, Douglas L., and Murali Chakinala. 2015. "Heart Failure: Pathophysiology and Diagnosis." In *Harrison's Principles of Internal Medicine,* edited by Dennis L. Kasper, Stephen L. Hauser, J. Larry Jameson, Anthony S. Fauci, Dan L. Longo, and Joseph Loscalzo. McGraw-Hill Education.

Marcellesi, Alexandre. 2015. "External Validity: Is There Still a Problem?" *Philosophy of Science* 82 (5): 1308–1317.

Marcum, James A. 2005. "Metaphysical Presuppositions and Scientific Practices: Reductionism and Organicism in Cancer Research." *International Studies in the Philosophy of Science* 19 (1): 31–45. https://doi.org/10.1080/02698590500051076.

Marks, Harry M. 1997. *The Progress of Experiment: Science and Therapeutic Reform in the United States, 1900–1990*. Cambridge University Press.

Marshall, M., J. Cornwell, and A. Collins. 2018. "Rethinking Medicine." *British Medical Journal* 363: k4987.

Martinez, Marco, and David Teira. 2024. "Why Experimental Balance Is Still a Reason to Randomize." *British Journal for the Philosophy of Science* 75 (2): 519–535. https://doi.org/10.1086/716096.

Massagué, Joan. 2008. "TGFβ in Cancer." *Cell* 134 (2): 215–230. https://doi.org/10.1016/j.cell.2008.07.001.

Mathie, Robert T., Suzanne M. Lloyd, Lynn A. Legg, et al. 2014. "Randomised Placebo-Controlled Trials of Individualised Homeopathic Treatment: Systematic Review and Meta-analysis." *Systematic Reviews* 3 (1): 142. https://doi.org/10.1186/2046-4053-3-142.

Mathie, Robert T., Nitish Ramparsad, Lynn A. Legg, Jürgen Clausen, Sian Moss, Jonathan R. T. Davidson, Claudia-Martina Messow, and Alex McConnachie. 2017. "Randomised, Double-Blind, Placebo-Controlled Trials of Non-individualised Homeopathic Treatment: Systematic Review and Meta-Analysis." *Systematic Reviews* 6 (1): 63. https://doi.org/10.1186/s13643-017-0445-3.

Matthews, J. Rosser. 1995. *Quantification and the Quest for Medical Certainty*. Princeton University Press.

Matthiessen, Dana. 2017. "Mechanistic Explanation in Systems Biology: Cellular Networks." *British Journal for the Philosophy of Science* 68 (1): 1–25. https://doi.org/10.1093/bjps/axv011.

Maung, H. H. 2016. "To What Do Psychiatric Diagnoses Refer? A Two-Dimensional Semantic Analysis of Diagnostic Terms." *Studies in History and Philosophy of Biological and Biomedical Sciences* 55: 1–10.

Mayo Clinic. 2024. "Statin Choice Decision Aid." Accessed July 1, 2024. https://statindecisionaid.mayoclinic.org/.

Mays, William H. 1880–1881. "On the supposed identity of the poisons of diphtheria, scarlatina, typhoid fever, and puerperal fever." *San Franciso Western Lancet* 9: 110–115.

McClimans, L. 2013. "The Role of Measurement in Establishing Evidence." *Journal of Medicine and Philosophy* 38 (5): 520–538.

McGarity, Thomas O., and Wendy E. Wagner. 2012. *Bending Science: How Special Interests Corrupt Public Health Research*. Harvard University Press.

McGinn, Colin. 1991. *The Problem of Consciousness*. Basil Blackwell.

McKeown, Thomas. 1976a. *The Modern Rise of Population*. Academic Press.

McKeown, Thomas. 1976b. *The Role of Medicine: Dream, Mirage, or Nemesis?* Nuffield Provincial Hospitals Trust.

McKitrick, Jennifer. 2003. "A Case for Extrinsic Dispositions." *Australasian Journal of Philosophy* 81: 155–174.

McMurray, J. J., S. Adamopoulos, S. D. Anker, et al. 2012. "ESC Guidelines for the Diagnosis and Treatment of Acute and Chronic Heart Failure 2012: The Task Force for the Diagnosis and Treatment of Acute and Chronic Heart Failure 2012 of the European Society of Cardiology. Developed in Collaboration with the Heart Failure Association (HFA) of the ESC." *European Heart Journal* 33 (14): 1787–1847. https://doi.org/10.1093/eurheartj/ehs104.

Mechtouff, L., L. Rascle, V. Crespy, E. Canet-Soulas, N. Nighoghossian, and A. Millon. 2021. "A Narrative Review of the Pathophysiology of Ischemic Stroke in Carotid Plaques: A Distinction Versus a Compromise Between Hemodynamic and Embolic Mechanism." *Annals of Translational Medicine* 9 (14): 1208. https://doi.org/10.21037/atm-20-7490.

Medical Research Council, Streptomycin in Tuberculosis Trials Committee. 1948. "Streptomycin Treatment for Pulmonary Tuberculosis." *British Medical Journal* 2: 769–782.

Medical Research Council, Whooping-Cough Immunization Committee. 1951. "The Prevention of Whooping Cough by Vaccination." *British Medical Journal* 1: 1463–1471.

Mercuri, Mathew. 2018. "Medical Nihilism." *Metapsychology Online Reviews* 22 (35).

Miles, Andrew. 2009. "On a Medicine of the Whole Person: Away from Scientistic Reductionism and Towards the Embrace of the Complex in Clinical Practice." *Journal of Evaluation in Clinical Practice* 15 (6): 941–949.

Mills, S. Anthony. 2023. "Why So Many Americans Are Losing Trust in Science." *New York Times*, August 3. https://www.nytimes.com/2023/10/03/opinion/science-americans-trust-covid.html.

Mitchell, Sandra D., and Angela M. Gronenborn. 2017. "After Fifty Years, Why Are Protein X-Ray Crystallographers Still in Business?" *British Journal for the Philosophy of Science* 68 (3): 703–723. https://doi.org/10.1093/bjps/axv051.

Molnar, George. 2006. *Powers: A Study in Metaphysics*. Oxford University Press.

Monroe, S. M., and K. L. Harkness. 2011. "Recurrence in Major Depression: A Conceptual Analysis." *Psychological Review* 118 (4): 655–674.

Monroe, Scott M., and Anne D. Simons. 1991. "Diathesis-Stress Theories in the Context of Life Stress Research: Implications for the Depressive Disorders." *Psychological Bulletin* 110 (3): 406–425.

Montano, C., T. Cassini, S. G. Ziegler, et al. 2022. "Diagnosis and Discovery: Insights from the NIH Undiagnosed Diseases Program." *Journal of Inherited Metabolic Disease* 45 (5): 907–918.

Morabia, Alfredo, ed. 2004. *A History of Epidemiologic Methods and Concepts*. Springer Online.

Morabia, Alfredo. 2006. "Pierre-Charles-Alexandre Louis and the Evaluation of Bloodletting." *Journal of the Royal Society of Medicine* 99: 158–160.

Morange, Michel. 1993. "The Discovery of Cellular Oncogenes." *History and Philosophy of the Life Sciences* 15 (1): 45–58. http://www.jstor.org/stable/23331545.

Morange, Michel. 1997. "From the Regulatory Vision of Cancer to the Oncogene Paradigm, 1975–1985." *Journal of the History of Biology* 30 (1): 1–29.

Morange, Michel. 1998. *A History of Molecular Biology*. Harvard University Press.

Morange, M. 2007. "The Field of Cancer Research: An Indicator of Present Transformations in Biology." *Oncogene* 26 (55): 7607–7610.

Morange, Michel. 2011. "History of Cancer Research." *Encyclopedia of Life Sciences*.

Morens, D. M., G. K. Folkers, and A. S. Fauci. 2008. "Emerging Infections: A Perpetual Challenge." *Lancet Infectious Diseases* 8 (11): 710–719.

Morgan, G. J. 2022. *Cancer Virus Hunters: A History of Tumor Virology*. Johns Hopkins University Press.

Moynihan, Ray, and Alan Cassels. 2006. *Selling Sickness: How the World's Biggest Pharmaceutical Companies Are Turning Us All into Patients*. Bold Type Books.

Moynihan, R., and R. Smith. 2002. "Too Much Medicine?" *British Medical Journal* 324 (7342): 859–860.

Mukerji, Nikil, and Edzard Ernst. 2022. "Why Homoeopathy Is Pseudoscience." *Synthese* 200 (5): 394.

Mukherjee, Siddhartha. 2010. *The Emperor of All Maladies: A Biography of Cancer*. Scribner.

Mulder, Renée L., Melissa M. Hudson, Smita Bhatia, et al. 2020. "Updated Breast Cancer Surveillance Recommendations for Female Survivors of Childhood, Adolescent, and Young Adult Cancer from the International Guideline Harmonization Group." *Journal of Clinical Oncology* 38 (35): 4194–4207. https://doi.org/10.1200/JCO.20.00562.

Mumford, Stephen. 2006. "The Ungrounded Argument." *Synthese* 149: 471–489.

Munson, R. 1981. "Why Medicine Cannot Be a Science." *Journal of Medicine and Philosophy* 6 (2): 183–208.

Murphy, Dominic. 2011. "Conceptual Foundations of Biological Psychiatry." In *Handbook of the Philosophy of Science*. Vol. 16: *Philosophy of Medicine*, edited by D. M. Gabbay, P. Thagard, and J. Woods. Elsevier.

Mushebenge, Aganze Gloire-Aimé, Samuel Chima Ugbaja, Nonkululeko Avril Mbatha, Rene B. Khan, and Hezekiel M. Kumalo. 2023. "A Comprehensive Analysis of Structural and Functional Changes Induced by SARS-CoV-2 Spike Protein Mutations" *COVID* 3 (9): 1454–1472.

Nagel, Ernest. 1961. *The Structure of Science: Problems in the Logic of Scientific Explanation*. Harcourt.

Nahin, R. L., A. Rhee, and B. Stussman. 2024. "Use of Complementary Health Approaches Overall and for Pain Management by US Adults." *JAMA* 331 (7): 613–615. https://doi.org/10.1001/jama.2023.26775.

Nasmith, L., P. Ballem, R. Baxter, et al. 2010. *Transforming Care for Canadians with Chronic Health Conditions: Put People First, Expect the Best, Manage for Results*. Canadian Academy of Health Sciences.

National Academies of Science, Engineering, and Medicine. 2016. *Biomarker Tests for Molecularly Targeted Therapies: Key to Unlocking Precision Medicine*. National Academies Press. https://doi.10.17226/21860.

National Academies of Science, Engineering, and Medicine. 2020. *Framework for Equitable Allocation of COVID-19 Vaccine*. National Academies Press. https://doi.org/10.17226/25917.

National Center for Complementary and Integrative Health. 2021a. "NCCIH Strategic Plan FY 2021–2025: Mapping the Pathway to Research on Whole Person Health." National Institutes of Health. https://nccih.nih.gov/about/nccih-strategic-plan-2021-2025.

National Center for Complementary and Integrative Health. 2021b. "Homeopathy: What You Need To Know." National Institutes of Health. Accessed March 9, 2024. https://www.nccih.nih.gov/health/homeopathy.

National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. 1978. *The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research*. US Department of Health, Education and Welfare.

National Health and Medical Research Council. 2015. "Evidence on the Effectiveness of Homeopathy for Treating Health Conditions." NHMRC Information Paper. Canberra.

National Health Service. 2024. "Predict Tool." University of Cambridge. https://breast.v3.predict.cam/tool.

National Institutes of Health Study Committee. 1965. *Biomedical Science and Its Administration: A Study of the National Institutes of Health. Report to the President.* The White House, Washington, DC.

National Research Council. 2011a. *Toward Precision Medicine: Building a Knowledge Network for Biomedical Research and a New Taxonomy of Disease*. National Academies Press.

National Research Council. 2011b. *Research Training in the Biomedical, Behavioral, and Clinical Research Sciences*. National Academies Press. https://www.ncbi.nlm.nih.gov/books/NBK56994/.

Ndumele, C. E., J. Rangaswami, S. L. Chow, et al. 2023. "Cardiovascular-Kidney-Metabolic Health: A Presidential Advisory from the American Heart Association." *Circulation* 148 (20): 1606–1635.

Neal, Jacob P. 2019. "When Causal Specificity Does Not Matter (Much): Insights from HIV Treatment." *Philosophy of Science* 86 (5): 836–846.

Nestsiarovich, Anastasiya, Jenna M. Reps, Michael E. Matheny, et al. 2021. "Predictors of Diagnostic Transition from Major Depressive Disorder to Bipolar Disorder: A Retrospective Observational Network Study." *Translational Psychiatry* 11 (1): 642.

Nissen, Steven E. 2010. "The Rise and Fall of Rosiglitazone." *European Heart Journal* 31: 773–776.

Nordenfelt, Lennart. 1995. *On the Nature of Health: An Action-Theory Approach*. 2nd ed. Kluwer.

Nordling, C. O. 1953. "A New Theory on Cancer-Inducing Mechanism." *British Journal of Cancer* 7 (1): 68–72.

Nordling, Linda. 2017. "Putting Genomes to Work in Africa." *Nature* 544: 20–22.

Norton, John D. 2021. *The Material Theory of Induction*. University of Calgary Press.

Nutton, Vivian. 2012. *Ancient Medicine*. Routledge.

Nyakutira, C., D. Röshammar, E. Chigutsa, et al. 2008. "High Prevalence of the CYP2B6 516G-->T(*6) Variant and Effect on the Population Pharmacokinetics of Efavirenz in HIV/AIDS Outpatients in Zimbabwe." *European Journal of Clinical Pharmacology* 64 (4): 357–365.

O'Connor, Cailin, and James Own Weatherall. 2020. "Hydroxychloroquine and the Political Polarization of Science." In *Thinking in a Pandemic: The Crisis of Science and Policy in the Age of COVID-19*, edited by Matt Lord. Boston Review & Verso Books. Original publication, https://www.bostonreview.net/articles/cailin-oconnor-james-owen-weatherall-covid-19-and-polarization/.

Ohshimo, Shinichiro, Josune Guzman, Ulrich Costabel, and Francesco Bonella. 2017. "Differential Diagnosis of Granulomatous Lung Disease: Clues and Pitfalls." Number 4 in the Series "Pathology for the Clinician," edited by Peter Dorfmüller and Alberto Cavazza 26 (145): 170012. https://doi.org/10.1183/16000617.0012-2017. https://err.ersjournals.com/content/errev/26/145/170012.full.pdf.

O'Leary, Diane. 2021. "Medicine's metaphysical morass: how confusion about dualism threatens public health." *Synthese* 199 (1–2): 1977–2005. https://doi.org/10.1007/s11229-020-02869-9.

Omran, Abdel R. 2005. "The Epidemiologic Transition: A Theory of the Epidemiology of Population Change, 1971." *Milbank Quarterly* 83 (4): 731–757. https://doi.org/10.1111/j.1468-0009.2005.00398.x.

Oreskes, Naomi. 2019. "Systematicity Is Necessary but Not Sufficient: On the Problem of Facsimile Science." *Synthese* 196 (3): 881–905.

Osler, William. 1907. "Introduction: The Evolution of Internal Medicine." In *Modern Medicine: Its Theory and Practice*, edited by William Osler and Thomas McCrae. Lea Brothers.

Osler, William. 1908. *An Alabama Student: And Other Biographical Essays*. Clarendon Press.

Osler, William. (1913) 2013. *The Evolution of Modern Medicine*: Project Gutenberg.

Osler, William, and John W. Churchman. 1907. "Syphilis." In *Modern Medicine: Its Theory and Practice*, edited by William Osler and Thomas McCrae. Lea Brothers.

Page, M. J., J. E. McKenzie, J. Kirkham, et al. 2014. "Bias Due to Selective Inclusion and Reporting of Outcomes and Analyses in Systematic Reviews of Randomised Trials of Healthcare Interventions." *Cochrane Database of Systematic Reviews* 2014 (10): Mr000035.

Pai, Madhukar, Marcel A. Behr, David Dowdy, et al. 2016. "Tuberculosis." *Nature Reviews Disease Primers* 2 (1): 16076. https://doi.org/10.1038/nrdp.2016.76.

Paneth, N., and M. J. Joyner. 2018. "Editors' Introduction to the Special Issue." *Perspectives in Biology and Medicine* 61 (4): 467–471.

Paneth, Nigel, Michael J. Joyner, and Arturo Casadevall. 2022. "Finding Evidence for Treatment Decisions in a Pandemic." *Trends in Molecular Medicine* 28 (7): 536–541.

Parascandola, Mark. 2005. "Epidemiology in Transition: Tobacco and Lung Cancer in the 1950s." In *Body Counts: Medical Quantification in Historical and Sociological Perspectives*, edited by Gerard Jorland, Annick Opinel, and George Weisz. McGill–Queen's University Press.

Parascandola, Mark. 2011a. "Causes, Risks, and Probabilities: Probabilistic Concepts of Causation in Chronic Disease Epidemiology." *Preventive Medicine* 53 (4–5): 232–234.

Parascandola, Mark. 2011b. "The Epidemiologic Transition and Changing Concepts of Causation and Causal Inference." *Revue d'histoire des sciences* 64 (2): 243–262.

Parascandola, M., and D. L. Weed. 2001. "Causation in Epidemiology." *Journal of Epidemiology and Community Health* 55 (12): 905–912. https://doi.org/10.1136/jech.55.12.905.

Park, Andrew, Daniel Steel, and Elicia Maine. 2023. "Evidence-Based Medicine and Mechanistic Evidence: The Case of the Failed Rollout of Efavirenz in Zimbabwe." *Journal of Medicine and Philosophy*. 48 (4): 348–358.

Park, J. J. H., G. Hsu, E. G. Siden, K. Thorlund, and E. J. Mills. 2020. "An Overview of Precision Oncology Basket and Umbrella Trials for Clinicians." *CA: A Cancer Journal for Clinicians* 70 (2): 125–137.

Parkkinen, V.-P., C. Wallmann, M. Wilde, et al. 2018. *Evaluating Evidence of Mechanisms in Medicine: Principles and Procedures*. Springer.

Pearce, N. 1996. "Traditional Epidemiology, Modern Epidemiology, and Public Health." *American Journal of Public Health* 86 (5): 678–683. https://doi.org/10.2105/ajph.86.5.678.

Pearl, Judea. 2009. *Causality: Models, Reasoning and Inference*. Cambridge University Press.

Pearl, Judea, and Elias Bareinboim. 2014. "External Validity: From Do-Calculus to Transportability Across Populations." *Statistical Science* 29 (4): 579–595.

Pearl, Judea, and Dana Mackenzie. 2018. *The Book of Why: The New Science of Cause and Effect*. Basic Books.

Pereira, Alexander. 2020. "Specific Phobia Is an Ideal Psychiatric Kind." *Philosophy, Psychiatry, & Psychology* 27 (3): 299–315. https://dx.doi.org/10.1353/ppp.2020.0037.

Pereira, T. V., R. I. Horwitz, and J. P. Ioannidis. 2012. "Empirical Evaluation of Very Large Treatment Effects of Medical Interventions." *Journal of the American Medical Association* 308 (16): 1676–1684.

Perillat, Lucie, and Mathew Mercuri. 2022. "Clinical Recommendations: The Role of Mechanisms in the GRADE Framework." *Studies in History and Philosophy of Science* 96: 1–9.

Petticrew, M., and I. Chalmers. 2011. "Use of Research Evidence in Practice." *Lancet* 378 (9804): 1696.

Pew Research Center. 2023. "Americans' Trust in Scientists, Positive Views of Science Continue to Decline." November. https://www.pewresearch.org/science/2023/11/14/confidence-in-scientists-medical-scientists-and-other-groups-and-institutions-in-society/.

Phillips, Christopher J. 2020. "Precision Medicine and Its Imprecise History." *Harvard Data Science Review* 2 (1): 1–10. https://doi.org/https://doi.org/10.1162/99608f92.3e85b56a.

Phillips, Christopher J. 2025. "Modeling the Epidemiologic Individual." *History of the Human Sciences*. https://doi.org/10.1177/09526951251337680.

Pierri, F., B. L. Perry, M. R. DeVerna, et al. 2022. "Online Misinformation Is Linked to Early COVID-19 Vaccination Hesitancy and Refusal." *Scientific Reports* 12 (1): 5966.

Pigliucci, Massimo. 2013. "The Demarcation Problem: A Belated Response to Laudan." In *Philosophy of Pseudoscience: Reconsidering the Demarcation Problem*, edited by Massimo Pigliucci and Maarten Boudry. University of Chicago Press.

Pigliucci, Massimo, and Maarten Boudry, eds. 2013. *Philosophy of Pseudoscience: Reconsidering the Demarcation Problem*. University of Chicago Press.

Plutynski, A. 2013. "Cancer and the Goals of Integration." *Studies in History and Philosophy of Biological and Biomedical Sciences* 44 (4 Pt. A): 466–476. https://doi.org/10.1016/j.shpsc.2013.03.019.

Plutynski, Anya. 2018. *Explaining Cancer: Finding Order in Disorder*. Oxford University Press.

Plutynski, Anya. 2020. "Cancer Modeling: The Advantages and Limitations of Multiple Perspectives." In *Understanding Perspectivism: Scientific Challenges and Methodological Prospects*, edited by Michela Massimi and Casey D. McCoy. Routledge.

Plutynski, Anya. 2021. "How Is Cancer Complex?" *European Journal for Philosophy of Science* 11 (2): 55.

Plutynski, Anya. 2022. "Why Precision Oncology Is Not Very Precise (and Why This Should Not Surprise Us)." In *Personalized Medicine in the Making*, edited by C. Beneduce and M. Bertolaso. Springer.

Polack, F. P., S. J. Thomas, N. Kitchin, et al. 2020. "Safety and Efficacy of the BNT162b2 mRNA Covid-19 Vaccine." *New England Journal of Medicine* 383 (27): 2603–2615. https://doi.org/10.1056/NEJMoa2034577.

Ponikowski, P., A. A. Voors, S. D. Anker, et al. 2016. "2016 ESC Guidelines for the Diagnosis and Treatment of Acute and Chronic Heart Failure: The Task Force for the Diagnosis and Treatment of Acute and Chronic Heart Failure of the European Society of Cardiology (ESC) Developed with the Special Contribution of the Heart Failure Association (HFA) of the ESC." *European Heart Journal* 37 (27): 2129–2200. https://doi.org/10.1093/eurheartj/ehw128.

Poole, C., I. Shrier, and T. J. VanderWeele. 2015. "Is the Risk Difference Really a More Heterogeneous Measure?" *Epidemiology* 26 (5): 714–718.

Popper, Karl R. 1963. *Conjectures and Refutations*. Routledge & Kegan Paul.

Porter, Roy. 1997. *The Greatest Benefit to Mankind: A Medical History of Humanity*. W. W. Norton.

Post, P. N., H. de Beer, and G. H. Guyatt. 2013. "How to Generalize Efficacy Results of Randomized Trials: Recommendations Based on a Systematic Review of Possible Approaches." *Journal of Evaluation in Clinical Practice* 19 (4): 638–643.

Powers, Alvin C. 2015a. "Diabetes Mellitus: Complications." In *Harrison's Principles of Internal Medicine*, edited by Dennis L. Kasper, Anthony S. Fauci, Stephen L. Hauser, Dan L. Longo, J. Larry Jameson and Joseph Loscalzo. McGraw-Hill Education.

Powers, Alan C. 2015b. "Diabetes Mellitus: Diagnosis, Classification, and Pathophysiology." In *Harrison's Principles of Internal Medicine*, edited by Dennis L. Kasper, Anthony S. Fauci, Stephen L. Hauser, Dan L. Longo, J. Larry Jameson and Joseph Loscalzo. McGraw-Hill Education.

Powers, Alvin C. 2015c. "Diabetes Mellitus: Management and Therapies." In *Harrison's Principles of Internal Medicine*, edited by Dennis L. Kasper, Anthony S. Fauci, Stephen L. Hauser, Dan L. Longo, J. Larry Jameson and Joseph Loscalzo. McGraw-Hill Education.

Prasad, Vinay. 2020. *Malignant: How Bad Policy and Bad Evidence Harm People with Cancer*. Johns Hopkins University Press.

Prasad, V., A. Vandross, C. Toomey, et al. 2013. "A Decade of Reversal: An Analysis of 146 Contradicted Medical Practices." *Mayo Clinic Proceedings* 88 (8): 790–798.

President's Council of Advisors on Science and Technology. 2008. "Priorities for Personalized Medicine." Office of Science and Technology Policy: Washington, D.C.

Pressler, T. R., and E. E. Kaizar. 2013. "The Use of Propensity Scores and Observational Data to Estimate Randomized Controlled Trial Generalizability Bias." *Statistics in Medicine* 32 (20): 3552–3568.

Prior, Elizabeth W., Robert Pargetter, and Frank Jackson. 1982. "Three Theses About Dispositions." *American Philosophical Quarterly* 19 (3): 251–257. http://www.jstor.org/stable/20013964.

Quazi, S. 2022. "Artificial Intelligence and Machine Learning in Precision and Genomic Medicine." *Medical Oncology* 39 (8): 120.

Quinones, Sam. 2015. *Dreamland: The True Tale of America's Opiate Epidemic*. Bloomsbury Press.

Quirke, Viviane, and Jean-Paul Gaudillière. 2008. "The Era of Biomedicine: Science, Medicine, and Public Health in Britain and France After the Second World War." *Medical History* 52 (4): 441–452.

Radden, Jennifer. 2003. "Is This Dame Melancholy? Equating Today's Depression and Past Melancholia." *Philosophy, Psychiatry, & Psychology* 10 (1): 37–52.

Rangachari, P. K. 1997. "Evidence-Based Medicine: Old French Wine with a New Canadian Label?" *Journal of the Royal Society of Medicine* 90 (5): 280–284.

Rather, L. J. 1959. "Towards a Philosophical Study of the Idea of Disease." In *The Historical Development of Physiological Thought*, edited by C. M. Brooks and P. F. Cranefield. Hafner.

Rather, L. J. 1978. *The Genesis of Cancer: A Study in the History of Ideas*. Johns Hopkins University Press.

RECOVER. 2024. "RECOVER: Researching COVID to Enhance Recovery." https://recovercovid.org/.

Reichenbach, Hans. 1949. *The Theory of Probability*. University of California Press.

Reilly, John J., Jr., Edwin K. Silverman, and Steven D. Shapiro. 2015. "Chronic Obstructive Pulmonary Disease." In *Harrison's Principles of Internal Medicine,* edited by Dennis L. Kasper, Stephen L. Hauser, J. Larry Jameson, Anthony S. Fauci, Dan L. Longo, and Joseph Loscalzo. McGraw-Hill Education.

Reiss, Julian. 2014. "What's Wrong with Our Theories of Evidence?" *Theoria: An International Journal for Theory, History and Fundations of Science* 29 (2): 283–306.

Reiss, Julian. 2015. "A Pragmatist Theory of Evidence." *Philosophy of Science* 82 (3): 341–362. https://doi.org/10.1086/681643.

Reiss, Julian. 2019. "Against External Validity." *Synthese* 196 (8): 3103–3121.

Relton, Clare, Katy Cooper, Petter Viksveen, et al. 2017. "Prevalence of Homeopathy Use by the General Population Worldwide: A Systematic Review." *Homeopathy* 106 (2): 69–78. https://doi.org/10.1016/j.homp.2017.03.002.

Rendle, Katharine A. S., Meghan C. Halley, Suepattra G. May, and Dominick L. Frosch. 2015. "Redefining Risk and Benefit: Understanding the Decision to Undergo Contralateral Prophylactic Mastectomy." *Qualitative Health Research* 25 (9): 1251–1259. https://doi.org/10.1177/1049732314557085.

Resnik, D. B., and K. C. Elliott. 2013. "Taking Financial Relationships into Account When Assessing Research." *Accountability in Research* 20 (3): 184–205.

Reutlinger, Alexander. 2020. "What Is Epistemically Wrong with Research Affected by Sponsorship Bias? The Evidential Account." *European Journal for Philosophy of Science* 10 (2): 15.

Reya, T., S. J. Morrison, M. F. Clarke, and I. L. Weissman. 2001. "Stem Cells, Cancer, and Cancer Stem Cells." *Nature* 414 (6859): 105–111. https://doi.org/10.1038/35102167.

Reznek, Lawrie. 1987. *The Nature of Disease*. Routledge & Kegan Paul.

Richmond, Phyllis A. 1980. "The Germ Theory of Disease." In *Times, Places, and Persons: Aspects of the History of Epidemiology*, edited by Abraham L. Lilienfeld. Johns Hopkins University Press.

Rochon, P. A., J. H. Gurwitz, R. W. Simms, et al. 1994. "A Study of Manufacturer-Supported Trials of Nonsteroidal Anti-Inflammatory Drugs in the Treatment of Arthritis." *Archives of Internal Medicine* 154 (2): 157–163.

Rose, Geoffrey. 2001. "Sick individuals and sick populations." *International Journal of Epidemiology* 30 (3): 427–32; discussion 433–434. https://doi.org/10.1093/ije/30.3.427.

Rose, Geoffrey. 2008. *The Strategy of Preventive Medicine*. Oxford University Press.

Rosenberg, C. E. 1977. "The Therapeutic Revolution: Medicine, Meaning, and Social Change in Nineteenth-Century America." *Perspectives in Biology and Medicine* 20 (4): 485–506.

Rosenberg, C. E. 2002. "The Tyranny of Diagnosis: Specific Entities and Individual Experience." *Milbank Quarterly* 80 (2): 237–260. https://doi.org/10.1111/1468-0009.t01-1-00003.

Ross, Lauren N. 2018. "The Doctrine of Specific Etiology." *Biology & Philosophy* 33 (5): 37. https://doi.org/10.1007/s10539-018-9647-x.

Ross, Lauren N. 2023. "Explanation in Contexts of Causal Complexity: Lessons from Psychiatric Genetics." In *Minnesota Studies in the Philosophy of Science. Vol. 23: From Biological Practice to Scientific Metaphysics*, edited by William C. Bausman, Janella K. Baxter, and Oliver M. Lean. University of Minnesota Press.

Ross, R. 1999. "Mechanisms of Disease—Atherosclerosis—An Inflammatory Disease." *New England Journal of Medicine* 340 (2): 115–126.

Roth, Phillip H. 2022. *Medicine as Science: The Making of Disciplinary Identity from Scientific Medicine to Biomedicine*. Nomos.

Rothman, Kenneth J. 1976. "Causes." *American Journal of Epidemiology* 104 (6): 587–592.

Rothman, Kenneth J. 1986. *Modern Epidemiology*. Little, Brown & Co.

Rothman, Kenneth J., and Sander Greenland. 2005. "Causation and Causal Inference in Epidemiology." *American Journal of Public Health* 95 (S1): S144–S150.

Rothman, Kenneth J., Sander Greenland, and Timothy L. Lash, eds. 2008. *Modern Epidemiology*. 3rd ed. Lippincott Williams & Wilkins.

Rothstein, William G. 1972. *American Physicians in the Nineteenth Century: From Sects to Science*. Johns Hopkins University Press.

Rothwell, P. M. 2005a. "External Validity of Randomised Controlled Trials: 'To Whom Do the Results of This Trial Apply?'" *Lancet* 365 (9453): 82–93.

Rothwell, P. M. 2005b. "Treating Individuals 2—Subgroup Analysis in Randomised Controlled Trials: Importance, Indications, and Interpretation." *Lancet* 365 (9454): 176–186.

Rothwell, P. M., M. Eliasziw, S. A. Gutnikov, et al. 2003. "Analysis of Pooled Data from the Randomised Controlled Trials of Endarterectomy for Symptomatic Carotid Stenosis." *Lancet* 361 (9352): 107–116.

Rothwell, P. M., M. Eliasziw, S. A. Gutnikov, C. P. Warlow, and H. J. Barnett. 2004. "Endarterectomy for Symptomatic Carotid Stenosis in Relation to Clinical Subgroups and Timing of Surgery." *Lancet* 363 (9413): 915–924.

Rothwell, P. M., Mehta, Z., Howard, S. C., Gutnikov, S. A., & Warlow, C. P. 2005. "Treating Individuals 3: From Subgroups to Individuals: General Principles and the Example of Carotid Endarterectomy." *Lancet* 365 (9455): 256–265.

Roush, Sherrilyn. 2009. "Second Guessing: A Self-Help Manual." *Episteme* 6 (3): 251–268.

Rush, Benjamin. 1805. *Medical Inquiries and Observations*. J. Conrad.

Sackett, David L., W. Scott Richardson, W. Rosenberg, and R. Brian Haynes. 1997. *Evidence-Based Medicine: How to Practice and Teach EBM*. Churchill Livingstone.

Sackett, D. L., W. M. C. Rosenberg, J. A. M. Gray, R. B. Haynes, and W. S. Richardson. 1996. "Evidence Based Medicine: What It Is and What It Isn't—It's About Integrating Individual Clinical Expertise and the Best External Evidence." *British Medical Journal* 312 (7023): 71–72.

Safer, D. 2002. "Design and Reporting Modifications in Industry-Sponsored Comparative Psychopharmacology Trials." *Journal of Nervous and Mental Disorder* 190: 583–592.

Savage, C. Wade. 1990. "Introduction." In *Minnesota Studies in the Philosophy of Science*. Vol. 14: *Scientific Theories*, edited by C. Wade Savage. University of Minnesota Press.

Schaffner, Kenneth F. 1980. "Theory Structure in the Biomedical Sciences." *Journal of Medicine and Philosophy* 5 (1): 57–97. https://doi.org/10.1093/jmp/5.1.57.

Schaffner, Kenneth F. 1993. *Discovery and Explanation in Biology and Medicine*. University of Chicago Press.

Schaffner, Kenneth F. 2006. "Reduction: The Cheshire Cat Problem and a Return to Roots." *Synthese* 151 (3): 377–402. https://doi.org/10.1007/s11229-006-9031-2.

Scheuermann, Richard H., Werner Ceusters, and Barry Smith. 2009. "Toward an Ontological Treatment of Disease and Diagnosis." *Summit on Translational Bioinformatics* 2009: 116–120.

Schleidgen, Sebastian, Corinna Klingler, Teresa Bertram, Wolf H. Rogowski, and Georg Marckmann. 2013. "What Is Personalized Medicine: Sharpening a Vague Term Based on a Systematic Literature Review." *BMC Medical Ethics* 14 (1): 55.

Schulson, Michael. 2024. "Are Evidence-Based Medicine and Public Health Incompatible?" Undark. https://undark.org/2024/02/21/evidence-based-medicine/.

Schwartz, P. H. 2008. "Risk and Disease." *Perspectives in Biology and Medicine* 51 (3): 320–334. https://doi.org/10.1353/pbm.0.0027.

Schwartz, P. H. 2011. "Questioning the Quantitative Imperative: Decision Aids, Prevention, and the Ethics of Disclosure." *Hastings Center Reports* 41 (2): 30–39.

Senn, Stephen. 2004. "Individual Response to Treatment: Is It a Valid Assumption?" *British Medical Journal* 329 (7472): 966–968.

Severinsen, M. 2001. "Principles Behind Definitions of Diseases—A Criticism of the Principle of Disease Mechanism and the Development of a Pragmatic Alternative." *Theoretical Medicine and Bioethics* 22 (4): 319–336. https://doi.org/10.1023/a:1011830602137.

Seyerle, A. A., and C. L. Avery. 2013. "Genetic Epidemiology: The Potential Benefits and Challenges of Using Genetic Information to Improve Human Health." *North Carolina Medical Journal* 74 (6): 505–508.

Shapiro, S. D. "Chronic Obstructive Pulmonary Disease." In *Harrison's Principles of Internal Medicine*, edited by Dan L. Longo. McGraw-Hill.

Shea, Beverley J., Jeremy M. Grimshaw, George A. Wells, et al. 2007. "Development of AMSTAR: A Measurement Tool to Assess the Methodological Quality of Systematic Reviews." *BMC Medical Research Methodology* 7 (1): 10.

Sheldon, T. A., G. H. Guyatt, and A. Haines. 1998. "Getting Research Findings into Practice—When to Act on the Evidence." *British Medical Journal* 317 (7151): 139–142.

Shortt, S. E. 1983. "Physicians, Science, and Status: Issues in the Professionalization of Anglo-American Medicine in the Nineteenth Century." *Medical History* 27 (1): 51–68.

Shrimanker, I., P. Tadi, and J. C. Sánchez-Manso. 2022. "Parkinsonism." StatPearls. https://www.ncbi.nlm.nih.gov/books/NBK542224/.

Shryock, Richard Harrison. 1974. *The Development of Modern Medicine: An Interpretation of the Social and Scientific Factors Involved*. University of Wisconsin Press.

Sidari, A., and E. Hill. 2018. "Diagnosis and Treatment of Gout and Pseudogout for Everyday Practice." *Primary Care* 45 (2): 213–236.

Siegel, R. L., K. D. Miller, H. E. Fuchs, and A. Jemal. 2021. "Cancer Statistics, 2021." *CA: A Cancer Journal for Clinicians* 71 (1): 7–33. https://doi.org/10.3322/caac.21654.

Simon, J. R. 2008. "Constructive Realism and Medicine—An Approach to Medical Ontology." *Perspectives in Biology and Medicine* 51 (3): 353–366.

Simon, J. R. 2010. "Advertisement for the Ontology for Medicine." *Theoretical Medicine and Bioethics* 31 (5): 333–346. https://doi.org/10.1007/s11017-010-9153-x.

Simon, Jeremy R. 2011. "Medical Ontology." In *Handbook of the Philosophy of Science*. Vol. 16: *Philosophy of Medicine*, edited by D. M. Gabbay, P. Thagard, and J. Woods. Elsevier.

Sismondo, S. 2008. "How Pharmaceutical Industry Funding Affects Trial Outcomes: Causal Structures and Responses." *Social Science & Medicine* 66 (9): 1909–1914.

Sliwa, P., and S. Horowitz. 2015. "Respecting All the Evidence." *Philosophical Studies* 172 (11): 2835–2858.

Smart, B. 2014. "On the Classification of Diseases." *Theoretical Medicine and Bioethics* 35 (4): 251–269. https://doi.org/10.1007/s11017-014-9301-9.

Smith, W. S., C. Johnston, and J. C. Hempill III. 2015. "Cerebrovascular Diseases." In *Harrison's Principles of Internal Medicine*, edited by Stephen L. Hauser, Dennis L. Kasper, J. Larry Jameson, Anthony S. Fauci, Dan L. Longo and Joseph Loscalzo. McGraw-Hill.

Sniderman, A. D., R. B. D'Agostino Sr., and M. J. Pencina. 2015. "The Role of Physicians in the Era of Predictive Analytics." *Journal of the American Medical Association* 314 (1): 25–26.

Snow, John. 1855. *On the Mode of Communication of Cholera*. London: John Churchill.

Snowden, F. M. 2008. "Emerging and Reemerging Diseases: A Historical Perspective." *Immunological Reviews* 225 (1): 9–26.

Snowden, Frank M. 2020. *Epidemics and Society: From the Black Death to the Present*. New Haven and London: Yale University Press.

Solomon, Miriam. 2015. *Making Medical Knowledge*. Oxford University Press.

Solomon, Miriam. 2020a. "After Disclosure." In *Uncertainty in Pharmacology*, edited by Adam La Caze and Barbara Osimani. Springer.

Solomon, Miriam. 2020b. "Medical Nihilism by Jacob Stegenga: After Medical Nihilism." *Studies in History and Philosophy of Biological and Biomedical Sciences* 81: 101271.

Solomon, Miriam, Jeremy Simon, and Harold Kinkaid, eds. 2017. *The Routledge Companion to Philosophy of Medicine*. Routledge.

Sonnenschein, Carlos, and Ana M. Soto. 1998. *The Society of Cells: Cancer and the Control of Cell Proliferation*. Routledge.

Sonnenschein, C., and A. M. Soto. 2000. "Somatic Mutation Theory of Carcinogenesis: Why It Should Be Dropped and Replaced." *Molecular Carcinogenesis* 29 (4): 205–211.

Spiegelhalter, David J. 2008. "Understanding Uncertainty." *Annals of Family Medicine* 6 (3): 196–197.

Spiegelman, D., and T. J. VanderWeele. 2017. "Evaluating Public Health Interventions: 6. Modeling Ratios or Differences? Let the Data Tell Us." *American Journal of Public Health* 107 (7): 1087–1091.

Sprenger, J., and J. Stegenga. 2017. "Three Arguments for Absolute Outcome Measures." *Philosophy of Science* 84 (5): 840–852.

Stacey, D., F. Légaré, K. Lewis, et al. 2017. "Decision Aids for People Facing Health Treatment or Screening Decisions." *Cochrane Database of Systematic Reviews* 4 (4): Cd001431.

Steel, Daniel P. 2008. *Across the Boundaries: Extrapolation in Biology and Social Science*. Oxford University Press.

Steel, Daniel. 2010. "A New Approach to Argument by Analogy: Extrapolation and Chain Graphs." *Philosophy of Science* 77 (5): 1058–1069.

Steel, Daniel. 2011. "Causal Inference and Medical Experiments." In *Handbook of the Philosophy of Science*. Vol. 16: *Philosophy of Medicine*, edited by D. M. Gabbay, P. Thagard, and J. Woods. Elsevier.

Steel, D. 2018. "If the Facts Were Not Untruths, Their Implications Were: Sponsorship Bias and Misleading Communication." *Kennedy Institute of Ethics Journal* 28 (2): 119–144.

Stegenga, Bauke T., Marjolein H. Kamphuis, Michael King, Irwin Nazareth, and Mirjam I. Geerlings. 2012. "The Natural Course and Outcome of Major Depressive Disorder in Primary Care: The PREDICT-NL Study." *Social Psychiatry and Psychiatric Epidemiology* 47 (1): 87–95.

Stegenga, J. 2014. "Down with the Hierarchies." *Topoi: An International Review of Philosophy* 33 (2): 313–322. https://doi.org/10.1007/s11245-013-9189-4.

Stegenga, J. 2015. "Measuring Effectiveness." *Studies in History and Philosophy of Biological and Biomedical Sciences* 54: 62–71. https://doi.org/10.1016/j.shpsc.2015.06.003.

Stegenga, Jacob. 2018. *Medical Nihilism*. Oxford University Press.

Stegenga, Jacob. 2020. "Medical Nihilism by Jacob Stegenga: Reply by the Author." *Studies in History and Philosophy of Biological and Biomedical Sciences* 81: 101274.

Stegenga, J. 2022. "Evidence of Effectiveness." *Studies in History and Philosophy of Science* 91: 288–295.

Stegenga, Jacob, and Tarun Menon. 2017. "Robustness and Independent Evidence." *Philosophy of Science* 84 (3): 414–435.

Stone, N. J., J. G. Robinson, A. H. Lichtenstein, et al. 2014. "2013 ACC/AHA Guideline on the Treatment of Blood Cholesterol to Reduce Atherosclerotic Cardiovascular Risk in Adults: A Report of the American College of Cardiology/American Heart Association Task Force on Practice Guidelines." *Circulation* 129: S1–45.

Straus, S. E., P. Glasziou, W. S. Richardson, and R. B. Haynes. 2011. *Evidence-Based Medicine: How to Practice and Teach It*. 4th ed. Churchill-Livingstone.

Strevens, Michael. 2007. "Mackie Remixed." In *Causation and Explanation*, edited by Joseph Keim Campbell, Michael O'Rourke, and Harry S. Silverstein. MIT Press.

Stuart, Elizabeth A., Stephen R. Cole, Catherine P. Bradshaw, and Philip J. Leaf. 2010. "The Use of Propensity Scores to Assess the Generalizability of Results from Randomized Trials." *Journal of the Royal Statistical Society Series A: Statistics in Society* 174 (2): 369–386.

Sturdy, Steve. 2011. "Looking for Trouble: Medical Science and Clinical Practice in the Historiography of Modern Medicine." *Social History of Medicine* 24 (3): 739–757.

Suárez, Mauricio. 2020. *Philosophy of Probability and Statistical Modelling: Elements in the Philosophy of Science*: Cambridge University Press.

Sung, N. S., W. F. Crowley Jr., M. Genel, et al. 2003. "Central Challenges Facing the National Clinical Research Enterprise." *Journal of the American Medical Association* 289 (10): 1278–1287.

Susser, M. 1985. "Epidemiology in the United States After World War II—The Evolution of Technique." *Epidemiologic Reviews* 7: 147–177.

Susser, M., and E. Susser. 1996. "Choosing a Future for Epidemiology. 1: Eras and Paradigms." *American Journal of Public Health* 86 (5): 668–673.

Susser, Mervyn. 1973. *Causal Thinking in the Health Sciences*. Oxford University Press.

Suzuki, E., and T. J. VanderWeele. 2018. "Mechanisms and Uncertainty in Randomized Controlled Trials: A Commentary on Deaton and Cartwright." *Social Science & Medicine* 210: 83–85. https://doi.org/10.1016/j.socscimed.2018.04.023.

Szasz, Thomas. 1961. *The Myth of Mental Illness*. Hoeber-Harper.

Szklo, Moyses, and F. Javier Nieto. 2007. *Epidemiology: Beyond the Basics*. 2nd ed. Jones and Bartlett.

Tabak, Lawrence. 2022. "How COVID-19 Immunity Holds Up over Time." *NIH Director's Blog* (blog), March 1.

Tabb, Kathryn. 2015. "Psychiatric Progress and the Assumption of Diagnostic Discrimination." *Philosophy of Science* 82 (5): 1047–1058.

Tabery, James. 2023. *Tyranny of the Gene: Personalized Medicine and Its Threat to Public Health*. Knopf.

Tabery, James. 2024. "'Precision Medicine' Is Genomic Medicine." *American Journal of Bioethics* 24 (3): 91–93.

Tan, Y. Y., V. Papez, W. H. Chang, S. H. Mueller, S. Denaxas, and A. G. Lai. 2022. "Comparing Clinical Trial Population Representativeness to Real-World Populations: An External Validity Analysis Encompassing 43 895 Trials and 5 685 738 Individuals Across 989 Unique Drugs and 286 Conditions in England." *Lancet Healthy Longevity* 3 (10): e674–e689.

Tartof, S. Y., J. M. Slezak, H. Fischer, et al. 2021. "Effectiveness of mRNA BNT162b2 COVID-19 Vaccine Up to 6 Months in a Large Integrated Health System in the USA: A Retrospective Cohort Study." *Lancet* 398 (10309): 1407–1416.

Taylor, F. Kraupl. 1979. *The Concept of Illness, Disease and Morbus*. Cambridge University Press.

Teerawattananon, Yot, Thunyarat Anothaisintawee, Chatkamol Pheerapanyawaranun, et al. 2022. "A Systematic Review of Methodological Approaches for Evaluating Real-World Effectiveness of COVID-19 Vaccines: Advising Resource-Constrained Settings." *PLoS One* 17 (1): e0261930.

Temkin, Oswei. 1977. "An Historical Analysis of the Concept of Infection." In *The Double Face of Janus and Other Essays in the History of Medicine*, edited by Oswei Temkin. Johns Hopkins University Press.

Teng, Shaolei, Anand K. Srivastava, Charles E. Schwartz, Emil Alexov, and Liangjiang Wang. 2010. "Structural Assessment of the Effects of Amino Acid Substitutions on Protein Stability and Protein Protein Interaction." *International Journal of Computational Biology and Drug Design* 3 (4): 334–349. https://doi.org/10.1504/IJCBDD.2010.038396.

Testa, J. R., M. Cheung, J. Pei, et al. 2011. "Germline BAP1 Mutations Predispose to Malignant Mesothelioma." *Nature Genetics* 43 (10): 1022–1025.

Tetlock, Philip E., and Dan Gardner. 2016. *Superforecasting: The Art and Science of Prediction*. Crown.

Thagard, Paul. 1999. *How Scientists Explain Disease*. Princeton University Press.

Thompson, R. Paul. 1983. *The Structure of Biological Theories*. State University of New York Press.

Thompson, R. Paul. 2011. "Theories and Models in Medicine." In *Handbook of the Philosophy of Science*. Vol. 16: *Philosophy of Medicine*, edited by D. M. Gabbay, P. Thagard, and J. Woods. Elsevier.

Tiller, E. C., N. B. Masters, K. L. Raines, et al. 2023. "Notes from the Field: Measles Outbreak—Central Ohio, 2022–2023." *Morbidity and Mortality Weekly Report* 72 (31): 847–849.

Timmermans, S., and M. Berg. 2003. *The Gold Standard: The Challenge of Evidence-Based Medicine and Standardization in Health Care*. Temple University Press.

Toews, I., A. Anglemyer, J. L. Nyirenda, et al. 2024. "Healthcare Outcomes Assessed with Observational Study Designs Compared with Those Assessed in Randomized Trials: A Meta-Epidemiological Study." *Cochrane Database of Systematic Reviews* 1 (1): Mr000034.

Tomes, Nancy. 1998. *The Gospel of Germs: Men, Women, and the Microbe in American Life*. Harvard University Press.

Tonelli, M. R., and B. H. Shirts. 2017. "Knowledge for Precision Medicine: Mechanistic Reasoning and Methodological Pluralism." *Journal of the American Medical Association* 318 (17): 1649–1650.

Tonelli, Mark, and Ross E. G. Upshur. 2019. "A Philosophical Approach to Addressing *Uncertainty* in Medical Education." *Academic Medicine* 94 (4): 507–511. doi:10.1097/ACM.0000000000002512.

Topol, Eric. 2019. *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again*. Basic Books.

Tresker, S. 2022. "Unreliable Threats: Conflicts of Interest Disclosure and the Safeguarding of Biomedical Knowledge." *Kennedy Institute of Ethics Journal* 32 (1): 103–126.

Triolo, V. A. 1964. "Nineteenth Century Foundations of Cancer Research: Origins of Experimental Research." *Cancer Research* 24: 4–27.

Tröhler, Ulrich. 2005. "Quantifying Experience and Beating Biases: A New Culture in Eighteenth-Century British Clinical Medicine." In *Body Counts: Medical Quantification in Historical and Sociological Perspectives*, edited by Gerard Jorland, Annick Opinel, and George Weisz. McGill–Queen's University Press.

Tröhler, U. 2011. "The Introduction of Numerical Methods to Assess the Effects of Medical Interventions During the 18th Century: A Brief History." *Journal of the Royal Society of Medicine* 104 (11): 465–474.

Truett, J., J. Cornfield, and W. Kannel. 1967. "A Multivariate Analysis of the Risk of Coronary Heart Disease in Framingham." *Journal of Chronic Diseases* 20 (7): 511–524.

Tsou, J. Y. 2011. "The Importance of History for Philosophy of Psychiatry: The Case of the DSM and Psychiatric Classification." *Journal of the Philosophy of History* 5 (3): 446–470.

Tsou, Jonathan Y. 2021. *Philosophy of Psychiatry*. Cambridge University Press.

Tsou, Jonathan Y. 2022. "Biological Essentialism, Projectable Human Kinds, and Psychiatric Classification." *Philosophy of Science* 89 (5): 1155–1165.

Turner, A. 2017. "Evaluating the UK House of Commons Science and Technology Committee's Position on the Implausible Effectiveness of Homeopathic Treatments." *Theoretical Medicine and Bioethics* 38 (4): 335–352.

Tuttle, Todd M., Elizabeth B. Habermann, Erin H. Grund, Todd J. Morris, and Beth A. Virnig. 2007. "Increasing Use of Contralateral Prophylactic Mastectomy for Breast Cancer Patients: A Trend Toward More Aggressive Surgical Treatment." *Journal of Clinical Oncology* 25 (33): 5203–5209. https://doi.org/10.1200/jco.2007.12.3141.

Uffelmann, Emil, Qin Qin Huang, Nchangwi Syntia Munung, et al. 2021. "Genome-Wide Association Studies." *Nature Reviews Methods Primers* 1 (1): 59.

Ullman, D. 2021. "An Analysis of Four Government-Funded Reviews of Research on Homeopathic Medicine." *Cureus* 13 (6): e15899. https://doi.org/10.7759/cureus.15899.

Upshaw, J. N., D. van Klaveren, M. A. Konstam, and D. M. Kent. 2018. "Digoxin Benefit Varies by Risk of Heart Failure Hospitalization: Applying the Tufts MC HF Risk Model." *American Journal of Medicine* 131 (6): 676–683.e2.

Upshur, Ross E. G. 2014. "Do Clinical Guidelines Still Make Sense? No." *Annals of Family Medicine* 12 (3): 202–203.

Upshur, Ross E. G., and Errol Colak. 2003. "Argumentation and Evidence." *Theoretical Medicine and Bioethics* 24: 283–299.

Upshur, Ross E. G., and Maya J. Goldenberg. 2020. "Countering Medical Nihilism by Reconnecting Facts and Values." *Studies in History and Philosophy of Science Part A* 84: 75–83.

Upshur, Ross E. G., Elizabeth G. VanDenKerkhof, and Vivek Goel. 2001. "Meaning and Measurement: An Inclusive Model of Evidence in Health Care." *Journal of Evaluation in Clinical Practice* 7 (2): 91–96.

Urbach, Peter. 1985. "Randomization and the Design of Experiments." *Philosophy of Science* 52 (2): 256–273. https://doi.org/10.1086/289243.

U.S. Surgeon General's Advisory Committee on Smoking and Health. 1964. *Smoking and Health: Report of the Advisory Committee to the Surgeon General of the Public Health Service.* U.S. Public Health Service.

VanderWeele, Tyler J., and James M. Robins. 2009. "Minimal Sufficient Causation and Directed Acyclic Graphs." *The Annals of Statistics* 37 (3): 1437–1465.

van Helvoort, Ton. 1999. "A Century of Research into the Cause of Cancer: Is the New Oncogene Paradigm Revolutionary?" *History and Philosophy of the Life Sciences* 21 (3): 293–330. http://www.jstor.org/stable/23332179.

Van Spall, Harriette G. C., Andrew Toren, Alex Kiss, and Robert A. Fowler. 2007. "Eligibility Criteria of Randomized Controlled Trials Published in High-Impact General Medical Journals: A Systematic Sampling Review." *Journal of the American Medical Association* 297 (11): 1233–1240.

Varga, Somogy. 2021. "Medicine as Science: Systematicity and Demarcation." *Synthese* 199 (1): 3783–3804.

Varga, Somogy. 2022. "The Criticism of Medicine at the End of Its 'Golden Age.'" *Theoretical Medicine and Bioethics* 43 (5–6): 401–419. https://doi.org/10.1007/s11017-022-09598-9.

Varga, Somogy. 2024. *Science, Medicine, and the Aims of Medicine: A Philosophical Analysis.* Cambridge University Press.

Vetter, B. 2014. "Dispositions without Conditionals." *Mind* 123: 129–156.

Virchow, Rudolf. (1847) 1956. "Standpoints in Scientific Medicine." *Bulletin of the History of Medicine* 30 (5): 436–449.

Virchow, Rudolf. 1855. "Cellular-Pathologie." *Archiv für pathologische Anatomie, Physiologie und klinische Medicin* 8: 3–39.

Viswanathan, M., S. Reddy, N. Berkman, et al. 2018. "Screening to Prevent Osteoporotic Fractures: Updated Evidence Report and Systematic Review for the US Preventive Services Task Force." *Journal of the American Medical Association* 319 (24): 2532–2551.

Vu, Thuy, and Francois X. Claret. 2012. "Trastuzumab: Updated Mechanisms of Action and Resistance in Breast Cancer." *Frontiers in Oncology* 2: 62–62. https://doi.org/10.3389/fonc.2012.00062.

Wagner, J. C., C. A. Sleggs, and P. Marchand. 1960. "Diffuse Pleural Mesothelioma and Asbestos Exposure in the North Western Cape Province." *British Journal of Industrial Medicine* 17 (4): 260–271.

Waks, A. G., and E. P. Winer. 2019. "Breast Cancer Treatment: A Review." *Journal of the American Medical Association* 321 (3): 288–300. https://doi.org/10.1001/jama.2018.19323.

Walker, M. J., J. Bourke, and K. Hutchison. 2019. "Evidence for Personalised Medicine: Mechanisms, Correlation, and New Kinds of Black Box." *Theoretical Medicine and Bioethics* 40 (2): 103–121.

Wang, H., and D. Yee. 2019. "I-SPY 2: A Neoadjuvant Adaptive Clinical Trial Designed to Improve Outcomes in High-Risk Breast Cancer." *Current Breast Cancer Reports* 11 (4): 303–310.

Wang, P. F., A. Neiner, and E. D. Kharasch. 2019. "Efavirenz Metabolism: Influence of Polymorphic CYP2B6 Variants and Stereochemistry." *Drug Metabolism and Disposition* 47 (10): 1195–1205.

Warner, John Harley. 1995. "The History of Science and the Sciences of Medicine." *Osiris* 10 (1): 164–193.

Warner, John Harley. 1997. *The Therapeutic Perspective: Medical Practice, Knowledge, and Identity in America, 1820–1885*. Princeton University Press.

Warner, John Harley. 1998. *Against the Spirit of System: The French Impulse in Nineteenth-Century American Medicine*. Princeton: Princeton University Press.

Watson, Thomas. 1858. *Lectures on the Principles and Practice of Physic*. Blanchard and Lea.

Watson, O. J., G. Barnsley, J. Toor, A. B. Hogan, P. Winskill, and A. C. Ghani. 2022. "Global Impact of the First Year of COVID-19 Vaccination: A Mathematical Modelling Study." *Lancet Infectious Diseases* 22 (9): 1293–1302.

Weber, Max. 1948. "Bureaucracy." In *Max Weber: Essays in Sociology*, edited by H. H. Gerth and C. Wright Mills. Routledge.

Weinberg, C. R. 2007. "Can DAGs Clarify Effect Modification?" *Epidemiology* 18 (5): 569–572. https://doi.org/10.1097/EDE.0b013e318126c11d.

Weinberg, Robert. A. 1988. "The Genetic Origins of Human Cancer." *Cancer* 61 (10): 1963–1968.

Weinberg, Robert A. 1999. *One Renegade Cell: How Cancer Begins*. Basic Books.

Weinberg, Robert A. 2014. *The Biology of Cancer*. 2nd ed. Garland Science.

Weinberger, A. H., M. Gbedemah, A. M. Martinez, D. Nash, S. Galea, and R. D. Goodwin. 2018. "Trends in Depression Prevalence in the USA from 2005 to 2015: Widening Disparities in Vulnerable Groups." *Psychological Medicine* 48 (8): 1308–1315.

Weindling, Paul. 1992. "From Infectious to Chronic Diseases: Changing Patterns of Sickness in the Nineteenth and Twentieth Centuries." In *Medicine in Society: Historical Essays,* edited by Andrew Wear, 303–316. Cambridge University Press.

Weinfurt, K. P., D. M. Seils, L. Lin, D. P. Sulmasy, A. B. Astrow, H. I. Hurwitz, R. B. Cohen, and N. J. Meropol. 2012. "Research Participants' High Expectations of Benefit in Early-phase Oncology Trials: Are We Asking the Right Question?" *Journal of Clinical Oncology* 30 (35): 4396–4400. https://doi.org/10.1200/jco.2011.40.6587.

Weiss, N. S. 2006. *Clinical Epidemiology: The Study of the Outcome of Illness*. Oxford University Press.

Weiss, P. 1977. "The System of Nature and the Nature of Systems: Empirical Holism and Practical Reductionism Harmonized." In *Toward a Man-Centered Medical Science*, edited by K. E. Schaefer, H. Hensel, and R. Brody. Future Press.

Weisz, George. 2005. "From Clinical Counting to Evidence-Based Medicine." In *Body Counts: Medical Quantification in Historical and Sociological Perspectives*, edited by Gerard Jorland, Annick Opinel, and George Weisz. McGill–Queen's University Press.

Weisz, George. 2014. *Chronic Disease in the Twentieth Century: A History*. Johns Hopkins University Press.

Weisz, G., A. Cambrosio, P. Keating, L. Knaapen, T. Schlich, and V. J. Tournay. 2007. "The Emergence of Clinical Practice Quidelines." *Milbank Quarterly* 85 (4): 691–727.

Welch, H. Gilbert, Lisa M. Schwartz, and Steve Woloshin. 2011. *Overdiagnosed: Making People Sick in the Pursuit of Health*. Basic Books.

Werkhoven, Sander. 2019. "A Dispositional Theory of Health." *British Journal for the Philosophy of Science* 70 (4): 927–952. https://doi.org/10.1093/bjps/axy003.

Westreich, D., J. K. Edwards, C. R. Lesko, E. Stuart, and S. R. Cole. 2017. "Transportability of Trial Results Using Inverse Odds of Sampling Weights." *American Journal of Epidemiology* 186 (8): 1010–1014. https://doi.org/10.1093/aje/kwx164.

Whelton, P. K., R. M. Carey, W. S. Aronow, et al. 2018. "2017 ACC/AHA/AAPA/ABC/ACPM/AGS/APhA/ASH/ASPC/NMA/PCNA Guideline for the Prevention, Detection, Evaluation, and Management of High Blood Pressure in Adults: Executive Summary: A Report of the American College of Cardiology/American Heart Association Task Force on Clinical Practice Guidelines." *Hypertension* 71 (6): 1269–1324. https://doi.org/10.1161/hyp.0000000000000066.

Whitbeck, C. 1977. "Causation in Medicine: The Disease Entity Model." *Philosophy of Science* 44 (4): 619–637. https://doi.org/10.1086/288771.

Whitehead, C. 2013. "Scientist or Science-Stuffed? Discourses of Science in North American Medical Education." *Medical Education* 47 (1): 26–32.

Wiegand, M., R. L. Eagan, R. Karimov, et al. 2023. "Global Declines in Vaccine Confidence from 2015 to 2022: A Largescale Retrospective Analysis." May 8. https://papers.ssrn.com/sol3/papers.cfm?abstract_id = 4438003.

Wilde, Michael, and Veli-Pekka Parkkinen. 2019. "Extrapolation and the Russo–Williamson Thesis." *Synthese* 196 (8): 3251–3262. https://doi.org/10.1007/s11229-017-1573-y.

Wilholt, Torsten. 2009. "Bias and Values in Scientific Research." *Studies in History and Philosophy of Science Part A* 40 (1): 92–101. https://doi.org/https://doi.org/10.1016/j.shpsa.2008.12.005.

Williams, N. E. 2007. "The Factory Model of Disease." *Monist* 90 (4): 555–584.

Williams, Neil E. 2011. "Arthritis and Nature's Joints." In *Carving Nature at Its Joints: Natural Kinds in Metaphysics and Science*, edited by Michael O'Rourke, Joseph Keim Campbell, and Matthew H. Slater. MIT Press.

Winokur, G., W. Coryell, M. Keller, J. Endicott, and H. Akiskal. 1993. "A Prospective Follow-Up of Patients with Bipolar and Primary Unipolar Affective Disorder." *Archives of General Psychiatry* 50 (6): 457–465.

Winslow, Charles Edward A. 1948. "What Is Epidemiology?" *American Journal of Public Health and the Nation's Health* 38: 854.

Witters, Dan. 2023. "U.S. Depression Rates Reach New Highs." Gallup, May 17. https://news.gallup.com/poll/505745/depression-rates-reach-new-highs.aspx.

Wolff, A. C., M. E. H. Hammond, K. H. Allison, et al. 2018. "Human Epidermal Growth Factor Receptor 2 Testing in Breast Cancer: American Society of Clinical Oncology/College of American Pathologists Clinical Practice Guideline Focused Update." *Journal of Clinical Oncology* 36 (20): 2105–2122. https://doi.org/10.1200/jco.2018.77.8738.

Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.
Woodward, James. 2010. "Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation." *Biology & Philosophy* 25 (3): 287–318.
Woodward, James, and Kenneth Kendler. 2023. "Polygene Risk Scores: A Philosophical Exploration." *Philosophy of Medicine* 4 (1): 1–21. https://doi.org/10.5195/pom.2023.156.
Wootton, David. 2007. *Bad Medicine: Doctors Doing Harm Since Hippocrates*. Oxford University Press.
World Health Organization. 2005. *Preparing a Health Care Workforce for the 21st Century: The Challenge of Chronic Conditions*. World Health Organization.
World Health Organization. 2014. *Global Status Report on Noncommunicable Diseases 2014*. World Health Organization.
World Health Organization. 2021. "COVID-19 Research and Innovation Achievements." World Health Organization.
Worrall, John. 2002. "What Evidence in Evidence-Based Medicine?" *Philosophy of Science* 69 (Proceedings): S316–S330.
Worrall, John. 2007. "Why There's No Cause to Randomize." *British Journal for the Philosophy of Science* 58 (3): 451–488. https://doi.org/10.1093/bjps/axm024.
Worrall, John. 2021. "Statins and CVD (Cardio-Vascular Disease): Now It's Personal!" LSE blog, June 29. https://www.lse.ac.uk/philosophy/blog/2021/06/29/statins-and-cvd-now-its-personal/.
Wright, Richard. 2011. "The NESS Account of Natural Causation: A Response to Criticisms." In *Perspectives on Causation*, edited by Richard Goldberg. Hart.
Xu, J., S. L. Murphy, K. D. Kochanek, et al. 2021. "Deaths: Final Data for 2019." *National Vital Statistics Reports* 70 (8): 1–86. https://www.cdc.gov/nchs/data/nvsr/nvsr70/nvsr70-08-508.pdf.
Yerushalmy, J., and C. E. Palmer. 1959. "On the Methodology of Investigations of Etiologic Factors in Chronic Diseases." *Journal of Chronic Diseases* 10 (1): 27–40.
Zachar, Peter. 2014. *A Metaphysics of Psychopathology*. MIT Press.
Zachar, Peter. 2019. "Metaphysical Problems in Psychiatric Classification and Nosology." In *The Bloomsbury Companion to Philosophy of Psychiatry*, edited by Serife Tekin and Robyn Bluhm. Bloomsbury Academic.
Zack, M. M., D. G. Moriarty, D. F. Stroup, et al. 2004. "Worsening Trends in Adult Health-related Quality of Life and Self-rated Health-United States, 1993-2001." *Public Health Reports* 119 (5): 493–505. https://doi.org/10.1016/j.phr.2004.07.007.
Zakrzewski, Wojciech, Maciej Dobrzyński, Maria Szymonowicz, and Zbigniew Rybak. 2019. "Stem Cells: Past, Present, and Future." *Stem Cell Research & Therapy* 10 (1): 68. https://doi.org/10.1186/s13287-019-1165-5.
Zamani, Mohammad, Shaghayegh Alizadeh-Tabari, Vahid Zamani, Javad Shokri-Shirvani, and Mohammad H. Derakhshan. 2022. "Worldwide and Regional Efficacy Estimates of First-Line Helicobacter pylori Treatments: A Systematic Review and Network Meta-Analysis." *Journal of Clinical Gastroenterology* 56 (2): 114–124.
Zanger, U. M., and M. Schwab. 2013. "Cytochrome P450 Enzymes in Drug Metabolism: Regulation of Gene Expression, Enzyme Activities, and Impact of Genetic Variation." *Pharmacology & Therapeutics* 138 (1): 103–141. https://doi.org/10.1016/j.pharmthera.2012.12.007.
Zarocostas, J. 2020. "How to Fight an Infodemic." *Lancet* 395 (10225): 676.
Zerhouni, E. 2003. "Medicine: The NIH Roadmap." *Science* 302 (5642): 63–72.
Zipkin, D. A., C. A. Umscheid, N. L. Keating, et al. 2014. "Evidence-Based Risk Communication: A Systematic Review." *Annals of Internal Medicine* 161 (4): 270–280.
Zwarenstein, Merrick, and Shaun Treweek. 2009. "What Kind of Randomized Trials Do We Need?" *Journal of Clinical Epidemiology* 62: 461–463.

# Index

*For the benefit of digital users, indexed terms that span two pages (e.g., 52–53) may, on occasion, appear on only one of those pages.*

*Note*: Tables and figures are indicated by an italic *t* or *f*.


absolute risk (AR), 233, 239–41, 244–45, 250–51, 282–83, 285. *See also* medical risk
  of heart attack, 252, 272, 276–77
  patient's risk *vs.,* 282–83, 284
  treated, 297–98
  untreated, 272, 276, 279, 297–98, 299
absolute risk reduction (ARR), 233, 290, 302, 373–77, 437, 440–41. *See also* relative risk reduction
  causal, 279–81, 282, 288, 300–4, 305–6
  expected utility difference, 374–75, 377
  measurement of, 284
  no non-targeted mechanisms, no paradoxical harms, 438
  no non-targeted mechanisms, paradoxical harms, 440
  non-targeted mechanisms, no paradoxical harms, 439
  preventive interventions with, 376
abstraction, 81, 187, 329
  unification by, 81, 100
Alzheimer's disease, 73–74, 200–1, 207, 209
American College of Cardiology/ American Heart Association (ACC/AHA), 239, 245, 250, 297, 298
Anjum, R., 254
antibiotics, 25, 48, 152, 171, 274–75, 372–73
antirealism. *See* fictionalism
antisepsis
  internal, 42, 400
  surgical, 22–23, 26–28, 72–73
argument framework, 257–58
argument theory of evidence, 16–17, 215, 218–19, 253–64, 294, 315, 362–63, 388, 391–92
Armitage, P., 85, 86–87
Aronowitz, R., 268
asthma, 175*t*, 178–79, 181–82
asymptomatic, 6–8, 10, 145–46, 171–72, 405
  chronic diseases, 166, 169, 189
  chronicity, asymptomaticity, and variable pathology, 10, 405
  dispositional, 198
  mental disorders, 163–64
  risk factors, 54, 211
  and risk reduction, 268–69, 403
atherosclerosis, 184, 271, 275, 276–77
autism spectrum disorder, 189

bacteriology, 33–34, 36–37, 39–40, 41–42, 72, 138–39, 149
  clinical medicine and, 47
  germ theory and, 83
  Louis Pasteur and, 26–27
Bareinboim, 312, 326–29
Bayesian/Bayesianism, 256n.28, 294–95, 380–81
Bayes's Rule, 367, 369, 378–80
Belmont Report, 35n.9
Bernard, C., 32–33, 40–41, 44, 223
  *Introduction to the Study of Experimental Medicine*, 40–41
  laboratory medicine, 44, 55, 223
biased evidence, 11, 16–17, 359, 360–61, 362–65, 392, 407. *See also* systemic bias
  and Medical Nihilism, 360–68
Bichat, X., 28–29
Binney, N., 180
biomarker(s), 275–76, 409–10, 411, 413, 414, 415–16
  biomarker-based clinical trials, 414, 416
  biomarker-based cohort studies, 416
biomedical model, 9, 13, 21, 23, 43, 67–68
Bird, A., 33
*Birth of the Clinic* (Foucault), 22–23
Blasimme, A., 106
Bliss, M., 2, 25
bloodletting, 25–26, 31–32, 36, 221, 222–23, 421
  Hippocratic humoral theory and, 35–36
  Louis's research on, 422, 424

Bluhm, R., 227, 254
Bolie, V. W., 79–80
Borgerson, K., 227, 254
breast cancer, 50–51, 277*f*, 277–78, 291–92, 335, 347–48, 375, 414. *See also* cancer
  BRCA1, 283–84
  HER2, 243–44, 347–48, 412
  PREDICT model, 243–44
  risk of, 283–84, 291–92
  trastuzumab for, 335, 347–48
British Doctors Study, 50–51
British Medical Research Council (MRC), 46–47, 50–51, 224, 400
Broadbent, A., 29–30, 133–34, 219–20, 280–81, 370–73, 386–87
  monocausal model, 122–23, 126–27

Cambrosio, A., 64–65, 243–44
cancer, 83–104, 114. *See also* breast cancer
  contagion and, 72–75
  driver genes in, 100–1, 114
  epidemics of, 48
  genetic mutation and, 86–87
  mortality *vs.* age, *85–86*
  multistage theory, 85–91
Cancer Hallmarks Hypothesis, 97, 99
Cancer Stem Cell (CSC) theory, 102–3, 109
carcinogenesis, 85, 87–88, 103. *See also* somatic mutation theory (SMT)
cardiovascular disease
  Framingham Heart Study, 132–33
  multifactorial thinking, 132
  risk factors for, 118, 132, 269
  statins and risk reduction, 269, 271–72
carotid artery stenosis, 278*f*, 335–36
carotid endarterectomy, 209–10, 275, 308. *See also* stroke
  effectiveness of, 336*f*
  stroke risk reduction and, 275, 278–79, 278*f*, 281, 331–32, 336*f*
Carter, K. C., 78–79, 107–8, 121–22, 144, 152–53
Cartwright, N., 254–56, 262, 312–13, 329–30, 391
categorical property, 179, 182, 183
causal graph, 129n.8, 321–22, 326, 329, 330–31
causal inferences, 49–50, 241, 246, 313, 352, 362–64, 388–89, 391–92
  biased, 362n.4, 363–64
  counterfactual approach to, 272–73
  effect size in RCT and, 234–35
  first-order therapeutic evidence supports, 392
  science of, science of extrapolation *vs.,* 312
  SSC theory and, 137
  theories of, 390
causal principles, 329, 330–31, 333
causation
  causal inference. *See* causal inferences
  multifactorial, 6–7, 129
  necessary, 123, 131
  probabilistic, 131–32, 137–38
  specific, 123, 131, 139, 144–45
  theory of, 88, 116, 136
Chalmers, I., 224–25
Chalmers, T., 66, 226
chance. *See* probability
Chin-Yee, B., 413
cholesterol, 54, 171–72, 259–60, 268–69, 275. *See also* statins
Cholesterol Treatment Trialists meta-analysis, 226–27
chronic diseases, 10, 11–12, 42n.16, 118, 163–64, 167–78
  asymptomatic, 171–78
  categorical base, 182–83
  causal base, 181–82, 183, 186
  characteristic manifestation, 185
  diabetes, 2–3, 169–70
  dispositional, 175*t*, 178–92
  epidemiology, 63n.30, 211
  HIV, 170
  nature of, 143, 163–93
  noncommunicable, 47–48, 73–74, 179
  risk factors for, 132, 133–34, 211, 267
  rival theories of, 186–92
  WHO's definition, 168
chronicity, 10, 163–64, 166, 167, 168, 171
  of chronic diseases, 167–71, 178
  chronicity, asymptomaticity, and variable pathology, 10, 405
  contemporary epidemic of, 211
chronic obstructive pulmonary disease (COPD), 141–43, 175*t*, 178–79, 181–82
classification, 29, 117, 120, 121, 140, 148, 161, 236–37
  causal, 125
  of communicable diseases, 117–18
  constitutive model of, 140–44
  of mental disorders, 193
  monocausal model of, 119, 121–29
  multifactorial model of, 117–18
  of noncommunicable diseases, 117–18, 140–44
  symptom-based, 207

clinical epidemiology, 53–54, 58–59, 66
contemporary, 33–34
evidence, 11, 53–54, 56–57, 64
numerical method *vs.,* 58–59
RCTs and, 52
clinical medicine, 33
diagnosis and prognosis, 29–30
with epidemiological science, 8
laboratory medicine model, 41–42
laboratory science and, 40, 41–42, 47
with modern research science, 15, 21, 33–34, 37, 39
orthodox. *See* orthodox medicine
Paris. *See* Paris clinical medicine
clinical priority, 21, 23–24, 44, 60–61, 67–68, 401
clinical trials, 28–29. *See* randomized controlled trial
biomarker-based, 414
Cholesterol Treatment Trialists meta-analysis of, 226–27
external validity of, 10, 308–11, 351–52, 407
extrapolation and, 309–11
generalizability limits, 308–9
patient-derived organoids, 411
precision medicine, 413–17
preventive and curative interventions, 271
clinicopathological correlation, 33–34, 37–38
Cochrane, A., 358–59, 361, 362
cohort study, 50–51, 55–56, 129–30, 284, 400–1, 432
biomarker-based, 416
Framingham Heart Study, 51, 132, 226
omics-centered, 413–14
prospective, 132, 168, 400–1
Collins, F., 408
comparative process tracing, 319–26
component causes, 135, 136, 150–51
concept of disease, 116, 148, 166–67, 173, 180, 414–15
constitutive model, 140–50
contemporary scientific medicine. *See* new modern medicine
Cooper, R., 204–6, 210–11
Cornfield, J., 66, 261–62
coronary artery disease (CAD), 140–41, 153–54, 171–72, 175*t*, 184–85. *See also* cardiovascular disease
counterfactual comparison, 272–73
COVID-19, 135, 249–50, 430–31, 432–33
hydroxychloroquine for, 429, 430–31, 432–33
ivermectin for, 429, 430–31
RCTs of vaccines, 249
resistance to scientific medicine, 17, 426–35
vaccination, 154, 308–9, 376, 429
vaccine skepticism, 419, 427, 429–30
credence, 270–71, 285, 294–99
benefits, 295–97
Devanasan's argument on, 299
independent evidence, 298
of patient probabilities, 295, 303–5
patient risk, 296, 300
propensity interpretation pitfalls and, 304
critical appraisal, 257, 263, 358–59, 391, 393
Cullen, W., 148–49
cure, 30, 169, 221, 273–74, 280–81, 301–2
effectiveness criterion, 26, 27–28, 30
interventions, 268, 270–75, 280–82, 305–6, 332

Dammann, O., 189
decision aid, 271–73. *See also* risk communication
degree of belief. *See* credence
Deductive-Nomological (DN) Model, 220–21
demarcation problems, 31, 32–33, 34, 36–37
dementia, 197*t*, 207, 209
depression, 193–94, 195–98, 197*t*, 200
Devanesan, A., 299, 370
diabetes mellitus (DM), type 1, 1, 2–3, 4–7, 178–79, 180–81
chronic disease, 169–70
dispositional theory of, 178–79, 181–82
multicausality, 118
subtypes, 207–8
diabetes mellitus (DM), type 2, 6–7, 169–70, 175*t*
insulin-responsive cells, 182, 207
management of, 7
*Diagnostic and Statistical Manual of Mental Disorders* (DSM), 192–200
diagnostic tests, 5, 22, 41, 56–57
Dickersin, K., 365
disease mechanisms, 88, 266–67
dispositional metaphysics of disease (DMD), 191
dispositions, 15–16, 176–77, 181–82, 183–84, 204–11
chronic diseases, 177–92
mental disorders, 192–203
disruption cofactors, 288, 332, 337
for atherosclerosis, 276–77
harm cofactors *vs.,* 279, 334–35
patient adherence and, 275–76
probability of, 438, 439, 440, 441
Dissemination Hypothesis, 77–78, 157

Distinguishability Hypothesis, 77–78, 125–26
Djulbegovic, B., 250
Doll, R., 85, 86–87
DSM. *See Diagnostic and Statistical Manual of Mental Disorders*
Dubos, R., 150

EBM. *See* evidence-based medicine
EBM movement, 61, 215–16, 218. *See also* evidence-based medicine
  clinical research evidence and, 226–27
  Cochrane's ideas, 358–59
  evidence hierarchies, 405–6
efavirenz, 248–49, 329–32
  metabolism by CYP2B6, 320–21, 326–28
  neuropsychiatric side effects, *320*, 343
*Effectiveness and Efficiency* (Cochrane), 358
effect size, 271, 281–82, 331–32, 333–34, 338
  extrapolation argument for, 333
  for preventive or curative medical intervention, 332
  problem of extrapolating, 329
  simple extrapolation, 340–41
Ehrlich, P., 152, 366
Elliott, K. C., 366
Engel, G., 9, 13, 64–65, 74–75, 110
epidemics, 1, 3, 47, 73–74
  of cancer and heart disease, 48
  cholera and smallpox epidemics, 25, 75–76, 427
  chronic diseases, 165
  emerging or re-emerging infections, 14, 48
  epidemiologic transitions, 47–54
  infectious disease
  mathematical theory of, 88n.13
  new diseases, 21, 23, 50, 73–74, 164–67
  noncommunicable disease, 6, 14
  epidemiological evidence, 3, 10, 16–17, 61, 66, 216, 217, 228
epidemiological medicine, 8, 9, 11, 14–15, 17, 21, 23–24, 55–68, 402–7
  earlier scientific medicines and, 57
  gold standard scientific method, 423
  interpretation, 424
  laboratory medicine *vs.,* 23–24, 59*t*
  modern, 434–35
  as precision medicine, 17
  special problems of, 407
  threefold epidemiologic transition, 403
epidemiological revolution, 400–2
epidemiological theory, 71–72, 74, 83–84, 85, 88, 89–91, 116, 140, 161–62
epidemiological uncertainty, 297–98, 304, 305
  credence and, 294–305
  patient probabilities as, 296
  prognostic and therapeutic evidence, 296–98
  theoretical relevance of, 304–5
*Epidemiologic Methods* (MacMahon, Pugh, and Ipsen), 134
epidemiologic transitions, 47–54. *See also* threefold epidemiologic transition
epidemiology, 49–54, 218, 225–26
  clinical. *See* clinical epidemiology
  contemporary scientific medicine, 60–61, 272–73, 399, 402, 407, 435
  external validity, 309–10, 351–54
  internal validity, 246–47
  in medical education, 401
  medical statistics and, 50n.23
  noncommunicable disease, 51, 52–53, 55, 56–57
  observational, 55–57
  risk and, 267
epilepsy, 148, 191
epistemic interpretations, 294–95
Erasmus, A., 365
Esterson, A., 201–2, 203
etiology, 184
  and causal classification, 125
  blueprint, 71–72, 74, 106–7, 108, 115
  definition, 123, 155–56
  of disease, 4–5
  monocausal, 127
  multifactorial, 10, 129–40
  multifactorial thinking, 129–40
  universal, 123–24
Evans, A., 152
evidence-based medicine (EBM), 8, 9, 16–17, 216–19, 223–24
  in academic medicine, 216–17
  authorities, 244–45
  evidence hierarchies, 10, 248–49, 253, 257, 262–64
  favoritism, 405–6
  GRADE approach, 425
  internal validity, 246–47
  measurement framework, 257–58, 264
  mission statement, 227
  movement, 61, 215–16, 218, 226–27
  Paris clinical medicine *vs.,* 58–59
  philosophical literature on, 217–18
  Risk GP Model, 215–16, 238–45, 264
  "risk-of-bias" tools, 383
  "rules of evidence" for medicine, 53–54
  traditionalists, 416
  Working Group, 283–84

*Evidence-Based Policy* (Cartwright and Hardie), 254, 309
evidence hierarchies, 10, 248–49, 253, 257, 262–64, 405–6
evidential problem of extrapolation, 314–15, 339–51. *See also* extrapolator's circle
  simple extrapolation, 339–45
  structural analogies, 345–51
evidence synthesis, 66, 257–58, 263–64, 263n.35
*The Evolution of Modern Medicine* (Osler), 23
expected utility (EU), 302
experimental medicine. *See* laboratory medicine
experimental physiology, 36–37, 38, 39, 42
experimental therapeutics, 42, 46
external validity, 310–14. *See also* extrapolation
  of clinical trials, 10, 308–11, 351–52, 407
  definition in clinical epidemiology, 310–12
extrapolation
  clinical trials relationship to, 309–11
  effect size and, 333, 340–41
  evidential problem of, 314–15, 339–51
  logical problem of, 319–38
  practical problems of, 351–54
  problems of. *See* problems of extrapolation
  simple, 314–15, 339–45
extrapolator's circle, 310, 316, 323

Feinstein, Alvan, 225–26
fictionalism
  about disease, 187
  about mental disorders, 201–2
first-order evidence, 388–90, 391–95
  effect estimates provided by, 394–95
  industry bias, 386–87
  publication bias, 386–87
Fisher, R. A., 51, 224, 261–62
Fleck, L., 43
Fletcher, R., 225–26
Fletcher, S., 225–26
Flexner, A., 45–47
Flexner Report, 45, 46–47
Food and Drug Administration (FDA), 61, 412
Foucault, M., 22–23, 28–29, 43
Framingham Heart Study, 66, 132–33, 168, 226, 268
Frank, S., 88–89

Gadsboll, N., 180
gastroesophageal reflux disease (GERD), 148, 171–72, 175*t*, 179, 181–82
Gaudilliere, J.-P., 67
generalizability. *See* external validity
general generalizability (GG), 341–42, 350–51. *See also* external validity
generalizability bias, 365
generalization, 245–50. *See also* extrapolation
genome-wide association studies (GWAS), 413–14, 415–16
Germ Hypothesis, 77–78, 107–8
germ theory, 27, 36–37, 41–42, 72–74, 75–83, 121
  central ideas of, 77
  Dissemination Hypothesis, 77–78
  Distinguishability Hypothesis, 77–79
  Germ Hypothesis, 77–78
  history of, 157
  monocausal model from, 152
  of putrefaction, 22–23, 42, 76
  theories of cancer in, 71
Gillies, D., 368–69
Glasziou, P., 243
Global Burden of Disease (GBD), 164–65
Goldacre, B., 361–62
Goodman, R. A., 174–76
Goodman, S., 283, 365
GRADE, 246–47. *See also* critical appraisal
Greene, J., 119–20, 268
Greenland, S., 233–34, 250
GWAS. *See* genome-wide association studies

Hahnemann, S., 36, 421–22
Hanahan, D., 93–94
Hardie, J., 254–56, 262, 312–13, 329–30
harm cofactors
  disruption cofactors *vs.,* 279, 334–35
  distribution of, 333
  intervention, interaction with, 332, 333
  operative stroke, 332
*Harrison's Principles of Internal Medicine* (Longo), 6–7, 149
Harvey, W., 34, 35–36
Hayward, R., 290
heart disease. *See* cardiovascular disease
hierarchy of evidence, 229–33, 237, 261–62, 264, 405–6
  EBM, 10, 53–54, 61, 64, 215–16, 217–18, 223–24, 227–29, 253, 257, 434
  problems with, 231–32, 237–38, 253, 263–64
  "quality of evidence," 235, 236, 246–47
  for treatment effectiveness, 230–31
higher-order evidence, 389–90
Hill, A. B., 55–56, 66, 224–25, 358, 400–1, 424
Hippocrates, 38
Hippocratic medicine, 26, 73, 84, 221, 222

history and philosophy of science (HPS), 12–13. *See also* philosophy of science
HIV disease, 146–47
Holmes, O. W., 422
homeopathy, 32–33, 36–37, 399, 420–26. *See also* pseudoscience
  classical, 420, 421–23, 425
  meta-analyses of RCTs, 402
Horwitz, A., 201–2, 203
hospital medicine. *See* Paris clinical medicine
Howick, J., 227
Hoyningen-Huene, P., 33
Hozo, Iztok, 250
human papillomavirus (HPV), 147
humorism, 26, 36–37, 39, 73
Hunter, J., 38
hypertension, 171–72, 175*t*, 185n.18

Illich, I., 357–58
inductive strategy, 174
internal validity, 246–47, 311–12, 364–65
industry bias, 364–65, 380, 381, 385–86, 393–95. *See also* systemic bias
  evidence of, 359, 381–82
  first-order evidence, 386–87
  interpretation, 394–95
  manifestation of, 383–84
  meta-evidentialist's approach, 383–84
  pharmaceutical, 359
  problem of, 384
infectious diseases, 42, 47, 50, 72, 80, 122, 123
  constitutive model, 143–44
  in developed countries, 47–48
  displaced deaths from, 90–91
  emerging or re-emerging, 14, 48
  genetics and immunity control, 128
  monocausal, 119–20, 121
  monocausal model *vs.* constitutive model, 144–50
  multifactorial etiology and, 138
  pathogenic cause of, 129–30
  research on, 71–72, 104–5, 106–7
  single-cause specificity, 128
  specific etiology, 117, 119, 127
  transition of, 49
  universal etiology for, 117–18, 143
inference theory of evidence. *See* argument theory of evidence
influenza, 149
  COVID-19 *vs.,* 154
Insel, T., 206–7
insulin, 2–3, 182, 207
International Agency for Research on Cancer, 319–20
*International Classification of Diseases* (ICD), 192–93
interventions, medical
  preventive and curative interventions, 268, 270–75, 280–82, 305–6, 332
  risk-reducing interventions, 268–69, 271
*Introduction to the Study of Experimental Medicine* (Bernard), 40–41
Ioannidis, J., 360–61
irrelevance objection, 390

Jenner, E., 26–27
Jewson, N. D., 43
Johns Hopkins medical school, 2, 34
Jolie, A., 283–84

Keating, P., 243–44
Kelly, R., 198
Kent, D., 290
Kermack, W. O., 79–80, 88n.13
King, L. S., 188
Klebs, E., 79
Knudson, A. G., Jr., 87, 88–89
Koch, R., 91–92, 122, 125, 400–1
Kravitz, R., 309–10, 343
Krieger, N., 72–89, 118
Kuhn, T., 43, 106

laboratory medicine, 5, 15, 21, 23, 41, 68
  biomedical model, 39–47
  epidemiological medicine *vs.,* 23–24, 59*t*
  equated with scientific medicine, 40
  experimental mindset, 39–40
  nineteenth-century schools, associated with, 43
Laennec, R., 148–49
Laing, R. D., 201–2, 203
Lakatos, I., 106, 107
Lange, M., 191
Laplane, L., 82, 102–3
Le Fanu, J., 25
Lemoine, M., 412
like comparison criterion, 420–21, 422–23, 424–25, 429. *See also* pseudoscience
  homeopathy and, 420
  Louis' description of, 421
  scientific medicine *vs.* unscientific medicine, 422
Lilienfeld, A., 261–62
Lister, J., 26–28, 42, 72–73, 76

logical problem of extrapolation, 319–38. *See also* extrapolator's circle
comparative process tracing, 319–26
mechanism-cofactor analysis, 329–38
transport formulas, 326–28
Louis, P. C. A., 27–28, 32–33, 40, 44, 58–59, 222–23, 421, 422, 424
Lundh, A., 383–84, 393–95
lung cancer, 85–86, 142. *See also* cancer
case-control study of, 400–1
epidemiological evidence on, 66
increasing incidence, 73–74
small-cell, 142–43, 157
smoking and, 50–51, 55–56, 85–86, 157, 160–61, 187–88, 261–62
younger *vs.* older age cohorts, 89–90

Mackie, J., 136–37
MacLagan, T. J., 2, 77, 78, 79, 193, 197*t*
manifest property, 179
Mant, D., 243
Marcellesi, A., 311–14
Marks, H., 225
Matthews, J. R., 44n.19, 225
Mays, W. H., 76–77, 79
McCrae, T., 2
McGinn, C., 205–6
McKendrick, A. G., 79–80, 88n.13
McKeown, T., 357–58
McMaster University in Canada, 53–54, 225–26
measurement, 262–63, 264
clinical outcomes, 215–16, 218, 228, 233, 238, 253, 257, 264, 267
epidemiological, 253, 254, 258, 262–63, 264
risk, 238, 240, 244
unbiasedness of, 263
measurement framework, 257–58
measurement theory, 215, 257
EBM's, 236–37, 264
of evidence, 218–19, 254, 258, 264
medical prediction, role in, 215–16
mechanisms, 63
biological, 65–66, 90, 237
biomedical, 221, 228
causal, 105–6, 108–10, 134
chancy targeted, 440
constitutive, 274–75
of disease, 88, 216
disruption, 274–75
etiological, 184, 186, 209, 274–75, 278–79, 331–32
genetic, 98
of homeopathic treatment, 424–25
mechanism-cofactor analysis, 329–38
molecular, 94, 108–9, 111, 216, 410
mutational, 92–93
no non-targeted, 438, 439
non-targeted, 335, 337–38, 438
pathological, 94, 270–71, 274–75, 281–82, 286, 287–88, 305–6
physiological, 216–17
pictorial and/or linguistic representations of, 80
specific, 108–9
targeted, 275–76, 332, 335–36, 337–38, 343–44
mechanism-cofactor analysis, 329–38. *See also* logical problem of extrapolation
medical education, 43–44, 45–46, 47
EBM and, 57
epidemiology in, 401
Flexner model, 45–47
modern American, 223
orthodox, 45
in Paris, 22–23, 44–45
physiology, 46
scientific, 44–45, 72–73
medical evidence, 11, 12, 13–14, 16–17, 57–58
argument theory of, 215, 218–19, 253–64
clinical epidemiology and, 54
EBM's framework for, 218
epidemiologic transition, 49
hierarchies of evidence, 10
misleading, 363
for predictions, 216–19
Medical Meta-Evidentialism *See* Meta-Evidentialism
medical model. *See* biomedical model; laboratory medicine
*Medical Nemesis* (Illich), 357–58
Medical Nihilism, 357–58, 359–60
biased evidence and, 360–68
broad-brush objection, 370–73
Meta-Evidentialism *vs.,* 356, 386
multiple measures of effectiveness objection, 373–77
objections, 356, 368–81
vague-probabilities objection, 377–81
medical prediction, 219–28
Deductive-Nomological Model, 220–21
EBM's framework for, 218
empiric models of, 221–22
empiric tradition, 215–16, 218–19, 221, 224, 238–39, 245, 264
mechanist tradition, 216–17, 219, 221, 223–24, 228, 237

Medical Nihilism (*cont.*)
medical evidence for, 216–19
prognosis, 216–17, 234, 241, 244–45, 251–52
Risk Generalization-Particularization Model, 215–16
medical risk, 267–71
absolute risk. *See* absolute risk; absolute risk reduction
epidemiology concept, 267
from individuals to populations, 271–82
from populations to individuals, 282–305
meaning of medical risk, 10, 267–71, 406
propensity interpretation of, 282–94
risk factors as disease targets, 267
relative risk. *See* relative risk; relative risk reduction
risk-reducing interventions, 268–69, 271
mental disorders, 163–64. *See also* depression
authoritative diagnostic resources for, 192–93
causal base, 200–1
characteristic syndrome of, 195–96
definition, 194
dispositional, 192–203, 197*t*
fictionalism about, 201
psychopathology of, 200–1, 203
syndrome-cause view, 195, 196–98
syndrome view, 195–98
Mercuri, M., 370
meta-analysis, 52, 242, 361, 366, 378, 381–82
Cholesterol Treatment Trialists, 226–27
of clinical trials, 308–9
EBM and, 405–6
of homeopathy RCTs, 402
industry sponsorship, 361, 384–85
medical interventions results, 388
of RCTs, 230, 238, 243–44, 258, 263–64, 271, 336*f*, 342, 369, 402, 405–6, 422–23, 434
of RRR, 242
trastuzumab, RCTs of, 243–44
meta-bias, 365
meta-evidence, 381–84, 391–92, 395–96
first-order evidence *vs.,* 388–90
Meta-Evidentialism, 16–17, 356, 359–60, 381–88, 391–92
adjustment procedure, 386–87
broad-brush objection, 386–87
effectiveness objection, 387
Medical Nihilism *vs.,* 356, 386
medical research problems, 386
meta-research findings, 393–95
methodological thesis, 383
"overestimation thesis," 382–83
traditional EBM and, 386
vague probabilities objection, 387–88
meta-research, 11, 381–82, 383–84, 388–89, 391
definition, 360–61
findings, 357, 359, 360, 364–65, 379–80, 382–83, 384, 386–87, 388–89, 393–96
by Lundh, 383–84
Meta-Evidentialism and, 387–88
meta-research evidence, 16–17, 359–60, 382–83. *See also* meta-evidence
argument from, 388–95
cognitive impact, 384–85
evidence of bias, 392
function of, 391–92
irrelevance objection, 390
screening-off objection, 393–94
underdetermination objection, 394
metastasis, 94, 97, 100
misleading evidence, 362–66, 392
indirect, 363
Meta-Evidentialism, 359–60
problem of, 356, 360, 362, 365, 366
mixed-bagism (therapeutic), 372–73
*Modern Epidemiology* (Rothman), 134, 135–36
modern medicine, 22–23, 24–31, 37, 68. *See also* new modern medicine
birth of, Wootton's argument on, 22–23, 25, 83
curative effectiveness of, 27, 28
development of, 24–25
making and remaking of, 1–4
modern scientific medicine as, 39
Osler's, 3, 8
science-based medicine, characterized by, 37
specific etiology, 15–16
therapeutic revolutions, 26
*Modern Medicine* (Osler & McCrae), 2, 4, 223
modern scientific medicine, 21, 31–39, 47, 55, 75–76, 222–23
demarcating from nonscientific or pseudoscientific medicine, 23–24
demarcation problems, 31, 32–33, 34, 36–37
era of, 8
in *The Evolution of Modern Medicine*, 23
numerical method and pathological anatomy, associated with, 33–34
pseudoscientific medicine *vs.,* 23–24
relevant sciences of, 39
as science-based medicine, 35–36
molecular biology, 65–66, 67, 110, 113, 114, 403–4, 410
genetic mutation, 74–75, 84
genetics or genomics, 67, 91, 92
laboratory research in, results of, 399, 418
sciences beyond, 67

theory-rich basic, 417
molecular epidemiology, 399, 403–4, 408, 414, 418
molecular diagnostics, 58–59, 412, 417
molecular reductionism, 9, 74–75, 110–15, 403–4
monocausal disease model, 117–18, 119–29
benefits of, 121
Carter's views on, 119–20
of classification, 121–29
constitutive model *vs.,* 144–50
deficiencies, 146
HIV disease, 146–47
myth of, 150–55
opportunistic infections, 146–47
prototypical diseases, 129–30
monocausal etiology, 127
monocausality, 118–21
monoclonal antibodies (mabs), 412
Morabia, A., 222–23
Morange, M., 92–93, 96, 106
multicausality, 118–21, 133
multifactorial etiology, 10, 129–40, 150–51, 404
multistage theory, 85–91

National Academies of Science, Engineering, and Medicine (NASEM), 409–10. *See also* National Research Council
National Center for Complementary and Integrative Health (NCCIH), 419, 425–26
National Institutes of Health (NIH), 46–47, 65–66
All of Us Precision Medicine Initiative, 413–14
Cancer Genome Atlas project, 413–14
National Cancer Institute, 66
Precision Medicine Initiative, 408, 432
RECOVER initiative, 432
National Research Council (NRC), 408, 409, 411–12, 413
Neal, N., 159–60
necessary cause, 123–24, 153–54, 156
*New England Journal of Medicine,* 47, 56
new medical model. *See* epidemiological medicine
new modern medicine, 1, 3, 8, 11, 12–13, 34–35, 39–40, 55, 63–64, 399
chronic disease in, 174
epidemiology, 60–61, 399, 402, 407, 435
epistemic movements in, 8
philosophical treatment of, 17
seven problems of, 9–12, 403
nihilism. *See* Medical nihilism
noncommunicable diseases, 48, 138–44, 164–65
constitutive model of classification, 140–44
epidemic, 6, 14
epidemiology, 138–39
morbidity and mortality from, 48, 49
multicomponent theory, 90
multifactorial etiology, 10
risk factors associated with, 139–40
singular etiology, 141
universal etiology and, 140
nonspecific pathology, 277–78
chronicity, asymptomaticity, and variable pathology, 10, 405
disease with, 158–59
specific pathology *vs.,* 209
Nordling, C. O., 85, 86–87
Norton, J., 256n.30
numerical method, 27–29, 36–37
clinical epidemiology *vs.,* 58–59
Louis's, 33, 40, 222–24, 357, 421
modern scientific medicine and, 33–34
Paris clinical medicine and, 34–35, 37–38, 39–40

Omran, A., 47–50, 90–91
oncogene paradigm. *See* somatic mutation theory (SMT)
Oncogene Hypothesis, 90, 98, 102
orthodox medicine, 8, 21, 22–23, 31–33
adopted science, 422
American, 44–45
naturalistic worldview, adopted, 38
Paris clinical medicine and, 44–45
premodern, 30
rationalistic medical systems, 36
Osler, W., 1–2, 3, 4–5, 8, 23, 34, 39, 223, 427
osteoporosis, 171–72, 175*t*, 178–79, 184–85

Paris clinical medicine, 27–29, 34–35, 37, 39–40, 43–45, 46
EBM *vs.,* 58–59
and laboratory medicine, 44, 60–61
model of, 37–38
practitioners, 47, 422
Park, A., 248–49
particularization, 250–53
Pasteur, L., 29, 76, 77–78, 91–92, 122
pathological anatomy, 28–29, 36–37, 38, 39, 42
clinical reasoning from, 37–38
epistemological primacy, 72
modern scientific medicine and, 33–34
and morbid anatomy, 33–34

pathological anatomy (*cont.*)
observational science of, 39–40
Paris clinical medicine and, 28–29, 34–35, 58–59
specific, 155–56
therapeutic progress facilitation, 29
tuberculosis as, 148–49
pathology, 184–85. *See also* specific pathology
pathophysiology, 36, 40–42, 185, 410
reasoning, 10, 53–54
patient-derived organoid (PDO), 411
Pearl, J., 309–10, 312, 326–29
pediatric oncology, 370, 376
pembrolizumab, 414
penicillin, 171
personalized or precision medicine, 17, 399, 401–2, 408–18
Chin-Yee argument on, 413
clinical trials, 413–17
conservative, 408
definition, 411–12
epistemology of, 416
exemplars of, 413
gene-centric, 67
patient-derived organoids, 411
radical, 409
research, 408
pharmaceutical revolution, 48, 51, 357
philosophy of medicine, 12–14
literature in, 32–33
physiological medicine. *See* laboratory medicine
physiology, 39–40, 41–42
digestive, 4–5
experimental, 33–34, 36–37, 38, 39, 42
pathogenesis and pathophysiology of diabetes, 6–7
pathological, 40–41
Pickering, G., 56
phenylketonuria, 191–92
Plutynski, A., 88, 106
polygenic risk scores, 139–40, 415–16
Post, P. N., 247, 339–40, 341, 342
postmodern medicine, 14–15, 17, 399, 402, 418–35
fall of scientific medicine, 426–35
rise of pseudoscientific medicine, 420–26
precision oncology, 74, 412–13
biomarker-based trials in, 414
therapeutic strategy in, 109
prediction. *See* medical prediction
preventive interventions, 268, 270–75, 280–82, 305–6, 332
preventive medicine, 267–68
probabilistic causation, 131–32
probability, 22, 251–53, 258, 276, 277, 286, 288–89, 292–95, 326–28, 329, 367–68, 424
calculus of, 58–59
of cardiovascular event, 261, 298
of causal variables and differences, 326–27
communicating, 304–5
epidemiological, 290–91, 292, 293–95
epistemic conceptions of, 284
epistemic interpretation, 294–95
frequentism, 285
interpretation of, 284, 285–86, 290–92, 294, 297–305
of outcome (clinical), 137–38, 217, 233, 240, 242–44, 250, 251, 273–74, 276, 321–22, 331–32, 337, 349–50
patient-level, 266–67, 270–71, 282–83, 287, 294, 296–97, 304, 305–6, 357
single-case, 250, 283, 285, 291, 293, 299
unbiased population assumption, 292–93
vague-probabilities objection, 377–81, 387–88
problems of extrapolation, 311–19
evidential, 314–15, 339–51
logical, 319–38
mechanisms approach to, 310
practical, 351–54
two, 311–19
process ontology of cancer, 143n.17
propensity interpretation, 270–71, 276, 298, 299–300, 304–5
distribution of individual propensities, *288–89*
effect size and, 299–300
epistemological and practical problems, 266–67
of patient-level probabilities, 292–93, 303
of patient's risk, 282–94
pitfalls of, 295–96, 304, 305–6
rejecting, 270–71
of single-case probabilities, 291
proximate causes, 148–49, 156–57
pseudoscience, 418–19, 426–27
during COVID-19 pandemic, 17, 420–26
homeopathy. *See* homeopathy
modern scientific medicine *vs.,* 23–24
rise of, 402, 420–26
publication bias, 359, 365, 385–86. *See also* systemic bias
first-order evidence, 386–87
graphical and statistical tests of, 393–94
industry bias and, 385–86
literature on, 361
meta-research evidence of, 389, 391–92

public health, 42n.17, 168
during COVID-19 pandemic, 402–34
epidemiology, 236, 261–62
medical care, impact of, 377
prevention of disease and, 267–68
universal etiology, 158

quantitative bias coefficient, 384–85
Quirke, V., 67

Radden, J., 195n.23
randomized controlled trial (RCT), 51–54, 217, 224–25, 253
bias risk, 424–25
COVID-19 pandemic, 249
difference-making evidence, 424–25
hierarchy's privileging of, 434
homeopathy, 402, 424–25
meta-analysis of, 230, 238, 258, 263–64, 271, 336*f*, 342, 369, 402, 405–6, 422–23, 434
rise of, 52
streptomycin for tuberculosis, 51–52, 400
rational credence, 294–96, 304. *See also* credence
realizable property. *See* disposition
reductionism, 74–75, 110–11
biomedical model, 113–14
molecular, 74–75, 110–15
relative risk reduction (RRR), 233, 302, 335, 437, 440–41. *See also* absolute risk reduction
causal, 279–81, 282, 288
distribution of, 290
measurement of, 284, 344
no non-targeted mechanisms, no paradoxical harms, 438
no non-targeted mechanisms, paradoxical harms, 440
non-targeted mechanisms, no paradoxical harms, 439
simple extrapolation, 340–41
Rendle, K., 291–92
Resnik, D. B., 366
respiratory infection, 146–47
*The Rise and Fall of Modern Medicine* (Le Fanu), 25
*The Rise of Causal Concepts of Disease* (Carter), 121–22
risk. *See* medical risk
risk communication, 283, 299–300, 304–5
Risk Generalization-Particularization (GP) Model, 215–16, 218–19, 238–45, 264, 363. *See also* generalization; particularization
argument theory of medical evidence, 253–64
for prognosis, 241
for statin effectiveness prediction, *258*
therapeutic predictions, 243–44
risk of generalizability biases, 365
risk reduction, 266–67, 270–71, 282, 300
absolute. *See* absolute risk reduction
ASCVD, 284
cardiovascular disease, 271–72
on credence interpretation, 300
diabetes management, contemporary, 3
effect size, 271
individual, 270, 275, 281, 288
interventions, 266–67, 270–71, 282, 300
multifactorial, 7
non-zero patient, 291
population-level, 284, 287, 289–90, 296
relative. *See* relative risk reduction
statin-mediated, 244–45
therapeutics of, 268–69
Rose, G., 88–90
Rosenberg, C., 26
Ross, L., 127–29, 133, 152–53
Rothman, K., 122, 233–34
Rous, P., 91–92
Rush, B., 36
Ryder, T., 2

Sackett, D., 225–26
SARS-CoV-2, 154, 249
Schaffner, K., 111
schizophrenia, 197*t*, 201–2, 203
science-based medicine, 31, 35–36, 37
scientific frameworks, 104, 106–7, 109–10, 116
germ theory and SMT as, 74, 83–84
scientific medicine, 15, 17, 30–39, 40, 45, 55, 59–60, 67–68, 75–76
biomedicine, 65, 67–68
during COVID-19 pandemic, 17, 426–35
definition, 8, 38
demarcation problems, 31, 32–33, 34
dominant model of, 15, 21, 60–61, 399, 402, 435
Flexner's vision, 45–46
Hill's model of, 56
model for physician training, 34
new model of, 21, 23–24
numerical method and pathological anatomy, 33–34
remodeling of, 57–58
transition in, 40, 44
scientific model, 64–65, 82, 104–5, 429

shared-cause specificity, 127
silent disorder phenomenon, 163–64, 171–78, 179, 181, 182, 212. *See also* asymptomatic
  chronic diseases and, 189
  DSM, 196
  interpretation of, 177–78
Simon, J., 173, 176
simple extrapolation, 314–15, 339–45
  effect size, 340–41
  evidential problem of extrapolation, 340–41
  general generalizability, 341–42
single-cause specificity, 127–29
Sismondo, S., 361–62
Smart, B., 147
smoking, 50–51, 55–56, 66, 85–86, 157
SMT. *See* somatic mutation theory
Snowden, F., 72–73
Solomon, M., 64, 223–24, 366, 370
Somatic Cell Hypothesis, 84, 96, 97–98, 99
somatic mutation theory (SMT), 71–72, 74–75, 90, 109, 115
  bacterial, 107
  cancer genes in, 87
  CSC theory, relationship with, 103
  dysregulated proliferation, 91
  multistage theory of cancer and, 84, 89–90
  reductionism critique, 112–15
  as scientific framework, 74
  theory of disease, 100
  unification of cancers, 100–2, 104
Sonnenschein, C., 96
Soto, A., 96
specific etiology, 121, 123
  and causal classification, 125
  doctrine of, Ross' views on, 127
  specific pathology *vs.,* 160–61
specific pathology, 117–18, 207–9
  antibiotic treatment of *M. tuberculosis* infection, 29
  chronic diseases and mental disorders associated with, 277–78
  dispositions and, 15–16, 204–11
  etiological factors, 155–61
  monocausal model, 15–16, 117, 121
  principle of, 155–56
  shared, 207–8
  specific etiology *vs.,* 160–61
Spiegelhalter, D., 305n.22
sponsorship bias. *See* industry bias
Sprenger, J., 302
statins, 268–69, 271, 275, 276–77, 279. *See also* cardiovascular disease
  decision aid, 271–73
  risk of heart attack, reduction in, *271*, 281
Steel, D., 311–12, 313, 316, 319–20, 321–22, 323–24, 325, 364
Stegenga, J., 302, 359–60, 366–69, 370–71, 372–75, 377–81, 382–83
stethoscope, 28–29, 30, 34–35
strep throat, 145–46, 147
streptomycin, 171. *See also* antibiotics
stroke, 118, 275, 278–79, 278*f*, 281, 331–32, 336*f*
structural analogies, 345–51
  evidential problem of extrapolation, 347–49
  structure-function relationships, 346, 350
  two factors function, 347–48
structural biology, 347
sufficient-component cause (SCC) model, 88n.13, 123–24, 134–38
  effect measures and, 437
  multifactorial etiology, 117–18
  with "potential outcomes" or "counterfactual framework," 437
  theory of causation, 136
Susser, M., 122, 126–27
Sydenham, T., 34, 188–89
syndrome view, 195–98
  -cause, 195, 196–98
syphilis, 42n.16, 171
systemic bias, 3, 11, 57–58, 381–88
  causal inference, 363–64
  industry. *See* industry bias
  inferential conception of, 363–64
  meta-bias, 365
  Meta-Evidentialism and, 16–17
  publication. *See* publication bias
Szasz, T., 201–2

targeted therapy, 109, 412
Taylor, F. K., 150
Thagard, P., 81n.10, 220
theory of disease, 104–5, 108–10
  dispositional theory, 186
  etiology, 117, 121, 133–34, 140, 161–62
  germ, 52–53, 72–73, 76
  prototypical modern, 75–83
  scientific, 73
  somatic mutation theory, 100, 104
therapeutic skepticism, 8
  contemporary, 359–60, 395–96
  cycles of, 357–60
Thompson, P., 79–80, 82
Thomsonianism, 31–32
threefold epidemiologic transition, 8, 15, 21, 23–24, 55–56
  in disease, evidence, and science, 49, 54, 68, 403

epidemiological medicine, rise of, 55
modern medicine upended by, 47
Tomes, N., 77
Toronto General Hospital, 4
transport formulas, 326–28
trastuzumab, for breast cancer, 243–44, 275–76
treatment effect heterogeneity, 266–67, 309
ARR aggregate, impact of, 305–6
individual, 277–79, 343
meta-analysis and, 342
trust, 431–32
tuberculosis, 42n.16, 48
pulmonary, 28–29, 51, 148–49
streptomycin for, 51–52, 400
tumors. *See also* cancer
CSC model, 103
heterotypic, 97–98
malignant, 100
microenvironment, 99
monoclonal, 97–98
tumorigenesis, 103
virus, 97
tumor suppressor gene
*APC* tumor suppressor gene, 100–1
causal role of, 108–9
mutations in, 94–96
*RB* tumor suppressor gene, 87, 92–93, 100–1, 112
type 1 diabetes mellitus. *See* diabetes mellitus (DM), type 1
type 2 diabetes mellitus. *See* diabetes mellitus (DM), type 2

umbrella trials, 414
underdetermination objection, 394
universal etiology, 123–24, 127, 138–39, 155, 187–88, 236–37. *See also* monocausal disease model
constitutive model for, 121
discovering, 157–58, 404
for infectious diseases, 117–18, 119, 143, 157, 209
of noncommunicable diseases, 140, 141
specific pathology as guide to, 155–61
target, 150–51
U.S. Preventive Services Task Force (USPSTF), 241–42, 245–46, 258
*Users' Guides to the Medical Literature*, 231–32, 251

vaccine skepticism, 419, 427, 429–30
vague-probabilities objection, 377–81, 387–88
Varga, S., 33
Varmus, H., 408
Virchow, R., 40–41, 55

Wakefield, J., 201–2, 203
Walker, M., 411n.2
Warner, J. H., 31–32, 40
Watson, T., 148–49
Weinberg, Robert, 92–94
Weisz, G., 166n.2
Werkhoven, S., 183–84
Whitbeck, C., 122, 126–27, 152–53, 185n.17
Williams, N., 191–92
Woodward, J., 125–26n.4, 159–60
Wooldridge, D. E., 65–66
Wooldridge Report, 65–66
Wootton, D., 22–23, 25, 26–30, 83
World Health Organization (WHO), 142, 168, 248–49

Zachar, P., 195n.23
zymotic theory, 75–76. *See also* germ theory