




Epidemiology

An Introduction



KENNETH J. ROTHMAN

Epidemiology

AN INTRODUCTION

KENNETH J. ROTHMAN

OXFORD
UNIVERSITY PRESS

2002

OXFORD
UNIVERSITY PRESS

Oxford New York
Athens Auckland Bangkok Bogotá Buenos Aires Cape Town
Chennai Dar es Salaam Delhi Florence Hong Kong Istanbul Karachi
Kolkata Kuala Lumpur Madrid Melbourne Mexico City Mumbai Nairobi
Paris São Paulo Shanghai Singapore Taipei Tokyo Toronto Warsaw

and associated companies in
Berlin Ibadan

Copyright © 2002 by Oxford University Press, Inc.

Published by Oxford University Press Inc.,
198 Madison Avenue, New York, New York 10016

Oxford is a registered trademark of Oxford University Press.

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
electronic, mechanical, photocopying, recording, or otherwise,
without the prior permission of Oxford University Press.

Library of Congress Cataloging-in-Publication Data

Rothman, Kenneth J.

Epidemiology : an introduction / Kenneth J. Rothman.

p. ; cm.

Includes bibliographical references and index.

ISBN 0-19-513553-9 (C : alk. paper) — ISBN 0-19-513554-7 (P)

1. Epidemiology. I. Title.

[DNLM: 1. Epidemiology. 2. Epidemiologic Methods.

WA 105 R846ea 2002] RA651 .R68 2002 614.4—dc21 2001036737

4 5 6 7 8 9

Printed in the United States of America
on acid-free paper.

Epidemiology

AN INTRODUCTION

To Emily, Margaret, and Samantha

Preface

Some observers of epidemiology appear to believe that epidemiology is little more than the application of statistical methods to the problems of disease occurrence and causation. In reality, epidemiology is much more than finely dressed statistics. It is a scientific discipline with roots in biology, logic, and the philosophy of science. For epidemiologists, statistical methods serve as an important tool, but not as a foundation. My aim in this book is to present a simple overview of the concepts that are the underpinnings of epidemiology, so that a coherent picture of epidemiologic thinking emerges for the student. The emphasis is not on statistics, formulas, or computation, but on epidemiologic principles and concepts.

For some, epidemiology is too simple to warrant serious attention, and for others it is too convoluted to understand. I hope to demonstrate to the reader that neither view is correct. The first chapter illustrates that epidemiology is more than just applying “common sense,” unless one has uncommonly good common sense. Although it is unusual to begin an epidemiology text with a discussion of confounding, I believe that the problem of confounding exemplifies why we need epidemiologic discipline to prevent our inferences from going astray. At the other extreme, those who believe that epidemiology is too complicated might think differently if they had a unifying set of ideas that extend across the boundaries of the many separate topics within epidemiology. My goal in this book has been to provide that unifying set of ideas.

These ideas begin with causation and causal inference, which are presented in Chapter 2. All too often these concepts are skipped over in scientific education. Nevertheless, for epidemiologists they are bedrock concerns that belong in any introduction to the field. Chapter 3 continues with a description of the basic epidemiologic measures, and Chapter 4 covers the main study types. An important thread for the student is the emphasis on measurement, and how to reduce or describe measurement error. Chapters 5 and 6 deal with measurement error. Systematic error, or bias, is treated first, in Chapter 5, and random error in Chapter 6. Chapter 7 introduces the basic analytic methods for estimat-

ing epidemiologic effects; these methods are extended in Chapter 8 to stratified data. Chapters 9 and 10 address the more advanced topics of interaction and multivariable modeling. These are subjects to be explored in more advanced courses, but their presentation here in elementary terms lays the groundwork for advanced study. It also draws a boundary between the epidemiologic approach to these topics and non-epidemiologic approaches that steer the analysis in the wrong direction. The final chapter deals with clinical epidemiology, a branch that is growing in scope and importance.

These topics are intended to constitute the core of a first course in epidemiology. Many epidemiology teachers will find that some subjects that they might like to include in such a course, such as the history of epidemiology, the study of infectious illness, or the social determinants of health and disease, have been omitted. To include these and other topics, however, would have made this a different book than the one I set out to write. My intent was not to create a comprehensive survey of the field, but rather a lean text that focuses on key conceptual issues.

Epidemiologic concepts are evolving, as any comparison of this text with earlier books will reveal. To complement the book, the publisher has graciously agreed to host a web site that will support reader participation in discussing, extending and revising points presented in the book. To begin, the web site will post contributed answers to the questions raised at the end of each chapter in the text. Interested readers can find the web site at <http://www.oup-usa.org/epi/rothman>.

Along the way I have received invaluable feedback from many students and colleagues. I am especially grateful to Kristin Anderson, Georgetowne Baghdady, Dan Brooks, Bob Green, Sander Greenland, Bettie Nelson, Ya-Fen Purvis, Igor Schillevoort, Bahi Takkouche, and Noel Weiss. Cristina Cann provided unflagging and generous encouragement. Katarina Augustsson deserves special mention for her careful reading of the manuscript and patient helpful criticisms. Finally, I am indebted to my colleague Janet Lang, who gently prodded me at the right time and was an inspiration throughout.

Contents

1. Introduction to Epidemiologic Thinking, 1
2. What Is Causation? 8
3. Measuring Disease Occurrence and Causal Effects, 24
4. Types of Epidemiologic Study, 57
5. Biases in Study Design, 94
6. Random Error and the Role of Statistics, 113
7. Analyzing Simple Epidemiologic Data, 130
8. Controlling Confounding by Stratifying Data, 144
9. Measuring Interactions, 168
10. Using Regression Models in Epidemiologic Analysis, 181
11. Epidemiology in Clinical Settings, 198
- Appendix, 218
- Index, 221

Introduction to Epidemiologic Thinking

This book presents the basic concepts and methods of epidemiology. Often considered the core science of public health, epidemiology involves “the study of the distribution and determinants of disease frequency”¹ or, put even more simply, “the study of the occurrence of illness.”²

The principles of epidemiologic research appear deceptively simple, which misleads some people into believing that anyone can master epidemiology just by applying common sense. The problem with this view is that the kind of common sense that is required may be elusive without training in epidemiologic concepts and methods. In this chapter, we glimpse some examples of the epidemiologic concept of confounding as a way to introduce epidemiologic thinking.

Common sense tells us that residents of Sweden, where the standard of living is generally high, should have lower death rates than residents of Panama, where poverty and more limited health care take their toll. Surprisingly, however, a greater proportion of Swedish residents than Panamanian residents die each year. This fact belies common sense. The explanation lies in the age distributions of the populations of Sweden and Panama. Figure 1–1 shows the population pyramids of the two countries. A *population pyramid* displays the age distribution of a population graphically. The population pyramid for Panama tapers dramatically from younger to older age groups, reflecting the fact that most Panamanians are in the younger age categories. In contrast, the population pyramid of Sweden is more rectangular, with roughly the same number of people in each of the age categories up to about age 60 and some tapering above that age. As these graphs make clear, Swedes tend to be older than Panamanians. For people of the same age in the two countries, the death rate among Swedes is indeed lower than that of Panamanians, but in both places older people die at a greater rate than younger people. Because Sweden has a population that is on the average older than that of Panama, a greater proportion of all Swedes die in

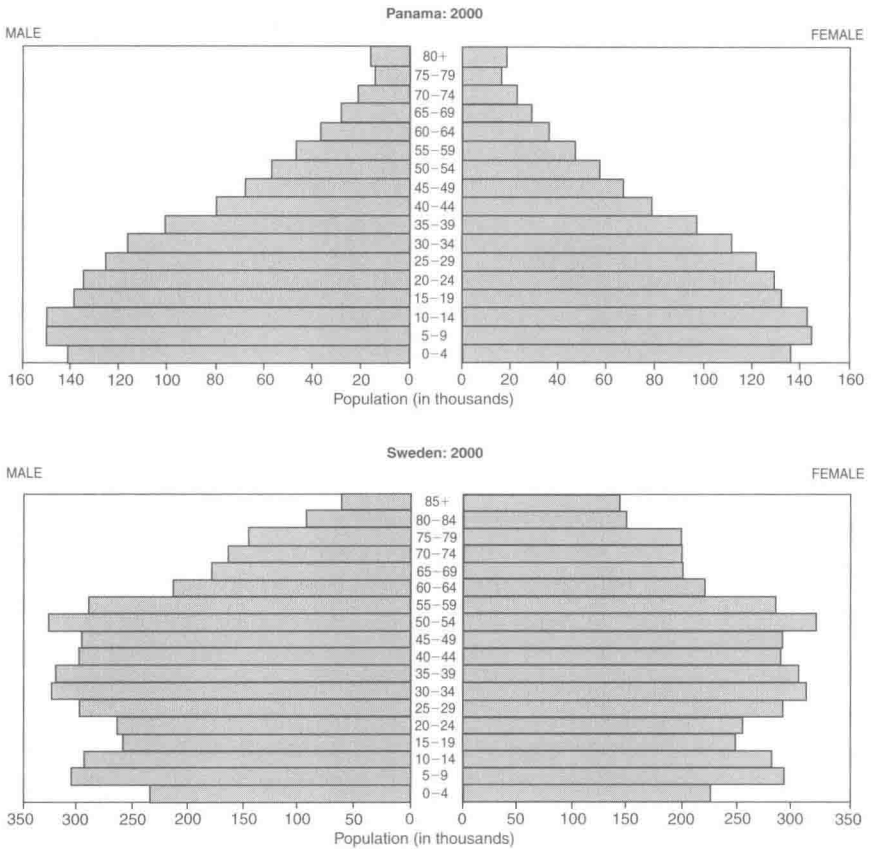


Figure 1-1. Age distribution of the populations of Panama and Sweden (population pyramids). Source: U.S. Census Bureau, International Data Base.

a given year, despite the lower death rates within age categories in Sweden compared with Panama.

This situation illustrates what epidemiologists call *confounding*. In this example, age differences between the countries are confounding the differences in death rates. Confounding occurs commonly in epidemiologic comparisons. Consider the following mortality data, summarized from a study that looked at smoking habits of residents of Whickham, England, in the period 1972–1974 and then tracked the survival over the next 20 years of those who were interviewed.³⁻⁵ Among 1314 women in the survey, nearly half were smokers. Oddly, proportionately fewer of the smokers died during the ensuing 20 years than nonsmokers. The data are reproduced in Table 1-1.

Only 24% of the women who were smokers at the time of the initial survey died during the 20-year follow-up period. In contrast, 31% of those who were nonsmokers died during the follow-up period. Does

Table 1–1. Risk of death in a 20-year period among women in Whickham, England, according to their smoking status at the beginning of the period*

Vital Status	Smoker	Nonsmoker	Total
Dead	139	230	369
Alive	443	502	945
Total	582	732	1314
Risk (dead/total)	0.24	0.31	0.28

*Data from Vanderpump et al.⁵

this difference indicate that women who were smokers fared better than women who were not smokers? Not necessarily. One difficulty that many readers quickly spot is that the smoking information was obtained only once, at the start of the follow-up period. Smoking habits for some women will have changed during the follow-up. Could those changes explain the results that appear to confer an advantage on the smokers? It is theoretically possible that all or many of the smokers quit soon after the survey and that many of the nonsmokers started smoking. While possible, this scenario is implausible, and without evidence for these changes in smoking behavior, this implausible scenario is not a reasonable criticism of the study findings. A more realistic explanation for the unusual finding becomes clear if we examine the data within age categories, as shown in Table 1–2 (the risks for each age group were calculated by dividing the number who died in each smoking group by the total of those dead or alive).

Table 1–1 combines all of the age categories listed in Table 1–2 into a single table, which is called the *crude* data. The more detailed display of the same data in Table 1–2 is called an *age-specific* display, or a display *stratified* by age. The age-specific data show that in the youngest and oldest age categories there was little difference between smokers and nonsmokers in risk of death. Few died among those in the younger age categories, regardless of whether they were smokers or not, whereas among the oldest women, nearly everyone died during the 20 years of follow-up. For women in the middle age categories, however, there was a consistently greater risk of death among smokers than nonsmokers, a pattern contrary to the impression gained from the crude data in Table 1–1.

Why did the nonsmokers have a higher risk of death in the study population as a whole? The reason is evident in Table 1–2: a much greater proportion of the nonsmoking women were in the highest age categories, the age categories that contributed a proportionately greater number of deaths. The difference in the age distributions between smokers and nonsmokers reflects the fact that, for most people, lifelong smoking habits are determined early in life. During the decades preced-

Table 1–2. Risk of death in a 20-year period among women in Whickham, England, according to their smoking status at the beginning of the period, by age*

Age (years)	Vital Status	Smoker	Nonsmoker	Total
18–24	Dead	2	1	3
	Alive	53	61	114
	Risk	0.04	0.02	0.03
25–34	Dead	3	5	8
	Alive	121	152	273
	Risk	0.02	0.03	0.03
35–44	Dead	14	7	21
	Alive	95	114	209
	Risk	0.13	0.06	0.09
45–54	Dead	27	12	39
	Alive	103	66	169
	Risk	0.21	0.15	0.19
55–64	Dead	51	40	91
	Alive	64	81	145
	Risk	0.44	0.33	0.39
65–74	Dead	29	101	130
	Alive	7	28	35
	Risk	0.81	0.78	0.79
75+	Dead	13	64	77
	Alive	0	0	0
	Risk	1.00	1.00	1.00

*Data from Vanderpump et al.⁵

ing the study in Whickham, there was a trend for increasing proportions of young women to become smokers. The oldest women in the Whickham study grew up during a period when few women became smokers, and they tended to remain nonsmokers for the duration of their lives. As time went by, a greater proportion of women who were passing through their teenage or young adult years became smokers. The result is a strikingly different age distribution for the female smokers and nonsmokers of Whickham. Were this difference in the age distribution ignored, one might conclude erroneously that smoking was not related to a higher risk of death. In fact, smoking *is* related to a higher risk of death, but confounding by age has obscured this relation in the crude data of Table 1–1. In Chapter 8, we return to these data and show how to calculate the effect of smoking on the risk of death after removing the age confounding.

Confounding is a problem that pervades many epidemiologic studies, but it is by no means the only issue that bedevils epidemiologic infer-

ences. One day, readers of the *Boston Globe*, a local newspaper, opened the paper to find a feature story about orchestra conductors. The point of the article was that conducting an orchestra was salubrious, as evinced by the fact that so many well-known orchestra conductors lived to be extremely old. Common sense suggests that if the people in an occupation tend to live long lives, the occupation must be good for health. Unfortunately, what appeared to be common sense for the author of the article is not very sensible from an epidemiologic point of view. The long-lived conductors cited in the article were mentioned because they lived to be old. Citing selected examples in this way constitutes *anecdotal* information, which can be extremely misleading. For all we know, the reporter searched specifically for examples of elderly conductors and overlooked other conductors who might have died at an earlier age. Most epidemiologists would not classify anecdotal information as epidemiologic data at all.

Furthermore, the reporter's observation has problems that go beyond the reliance on anecdotes instead of a formal evaluation. Suppose that the reporter had identified all orchestra conductors who worked in the United States during the past 100 years and studied their longevity. This approach would avoid relying on hand-picked examples, but it still suffers from an important problem that would lead to an incorrect answer. The problem is that orchestra conductors are not born as orchestra conductors. They become conductors at a point in their careers when they may have already attained a respectable age. If we start with a group of people who are 40 years old, on the average they are likely to survive to an older age than the typical person who was just born. Why? Because they have a 40-year head start; if they died before age 40, they could not be part of a group in which everyone is 40 years old. To find out if conducting an orchestra is beneficial to health, we should compare the risk of death among orchestra conductors with the risk of death among other people who have attained the same age as the conductors. Simply noting the average age at death of the conductors will give the wrong answer, even if all orchestra conductors were to be studied.

Here is another example that makes this point clearly. Suppose that we study two groups of people and look at the average age at death among those who die. In group A, the average age at death is 4 years; in group B, it is 28 years. Can we say that being a member of group A is riskier than being a member of group B? We cannot, for the same reason that the age at death of orchestra conductors was misleading. Suppose that group A comprises nursery school students and group B comprises military commandos. It would be no surprise that the average age at death of people who are currently military commandos is 28 years or that the average age at death of people who are currently nursery school students is 4 years. Still, we suspect that being a military commando is riskier than being a nursery school student and that these data on the

average age at death do not address the issue of which of these groups faces a greater risk of death. When one looks at the average age at death, one looks only at those who actually die and ignores all those who survive. Consequently, average age at death does not reflect the risk of death but only a characteristic of those who die.

In a study of factory workers, an investigator inferred that the factory work was dangerous because the average age at onset of a particular kind of cancer was lower in these workers than among the general population. But just as for the nursery school students and the military commandos, if these workers were young, the cancers that occurred among them would have to be occurring in young people. Furthermore, the age at onset of a disease does not take into account what proportion of people get the disease.

These examples reflect the fallacy of comparing the average age at which death or disease strikes rather than comparing the risk of death between groups of the same age. We will explore the proper way to make epidemiologic comparisons in later chapters. The point of these examples is to illustrate that a common-sense approach to a simple problem can be overtly wrong, until we educate our common sense to appreciate better the nature of the problem. Any sensible person can understand epidemiology, but without considering the principles outlined in this book, even a sensible person using very common sense is apt to go astray. By mastering a few fundamental epidemiologic principles, it is possible to refine our common sense to avoid these traps.

Questions

1. Age is a variable that is often responsible for confounding in epidemiology, in part because the occurrence of many diseases changes with age. The change in disease risk with age is often referred to as the effect of age. Does it make sense to think of age as having an effect on disease risk, or is it more sensible to think that the effect of age is itself confounded by other factors?
2. More people in Los Angeles die from cardiovascular disease each year than do people in San Francisco. What is the most important explanation for this difference? What additional factors would you consider to explain the difference in the number of deaths?
3. In Table 1–2, which age group would you say shows the greatest effect of smoking on the risk of death during the 20-year interval? How have you defined “greatest effect”? What other way could you have defined it? Does your answer depend on which definition you use?
4. On a piece of graph paper, use the data in Table 1–2 to plot the 20-year risk of death against age. Put age on the horizontal axis and the 20-year risk of death on the vertical axis. Describe the shape of the curve. What biologic forces account for the shape?

5. A physician who was interested in jazz studied the age at death of jazz musicians, whom he identified from an encyclopedia of jazz. He found that the average age at death of the jazz musicians was about the same as that of the general population. He concluded that this finding refuted the prevailing wisdom that jazz musicians tended to live dissolute lives and thus experienced greater mortality than other people. Explain his error.
6. A researcher determined that being left-handed was dangerous because he found that the average age at death of left-handers was lower than that of right-handers. Was he correct? Why or why not?
7. What is the underlying problem in comparing the average age at death or the average age at which a person gets a specific disease between two populations? How should you avert this problem?

References

1. MacMahon, B, Pugh, TF: *Epidemiology: Principles and Methods*, Chapter 1, Boston: Little, Brown, 1970.
2. Cole, P: The evolving case-control study. *J Chron Dis* 1979;32:15–27.
3. Appleton, DR, French, JM, Vanderpump, MPJ: Ignoring a covariate: an example of Simpson's paradox. *Am Statistician* 1996;50:340–341.
4. Tunbridge, WMG, Evered, DC, Hall, R, et al.: The spectrum of thyroid disease in a community. *Clin Endocrinol* 1977;7:481–493.
5. Vanderpump, MPJ, Tunbridge, WMG, French, JM, et al: The incidence of thyroid disorders in the community: a twenty-year follow-up of the Whickham survey. *Clin Endocrinol* 1995;43:55–69.

2

What Is Causation?

The acquired wisdom that certain conditions or events bring about other conditions or events is an important survival trait. Consider an infant whose first experiences are a jumble of sensations that include hunger pangs, thirst, color, light, heat, cold, and many other stimuli. Gradually, the infant begins to perceive patterns in the jumble and to anticipate connections between actions such as crying and effects such as being fed. Eventually, the infant assembles an inventory of associated perceptions. We can imagine that the concept slowly develops that some of these phenomena are causally related to others that follow. Along with this growing appreciation for specific causal relations comes the general concept that some events or conditions can be considered causes of other events or conditions.

Thus, our first appreciation of the concept of causation is based on our own observations. These observations typically involve causes with effects that are immediately apparent. For example, when one changes the position of a light switch on the wall, one can see the instant effect of the light going on or off. There is more to the causal mechanism for getting the light to shine than turning the light switch to the "on" position, however. Suppose the electric lines to the building are down in a storm. Turning on the switch will have no effect. Suppose the bulb is burned out. Again, the switch will have no effect. One cause of the light going on is having the switch in the proper place, but along with it we must include a supply of power to the circuit, a working bulb, and wiring. When all other factors are already in place, turning the switch will cause the light to go on, but if one or more of the other factors is not playing its causal role, the light will not go on when the switch is turned. There is a tendency to consider the switch to be the unique cause of turning on the light, but in reality we can define a more intricate causal mechanism, in which the switch is one component of several. The tendency to identify the switch as the unique cause stems from its usual role as the final factor that acts in the causal mechanism. The wiring can be considered part of the causal mechanism, but once it is put in place, it seldom warrants further attention. The switch, however, is often

the only part of the mechanism that needs to be activated to obtain the effect of turning on the light. The effect usually occurs immediately after turning the switch, and as a result we slip into a frame of thinking in which we identify the switch as a unique cause. The inadequacy of this assumption is emphasized when the bulb goes bad and needs to be replaced.

The Causal Pie Model

Causes of disease can be conceptualized in the same way as the causes of turning on a light. A helpful way to think about causal mechanisms of disease is depicted in Figure 2–1.¹ Each “pie” in the diagram represents a theoretical *causal mechanism* for a given disease, sometimes called a “sufficient cause.” There are three pies, to illustrate that there are multiple mechanisms that cause any type of disease. Each individual instance of disease will occur through a single mechanism or sufficient cause. A given causal mechanism requires the joint action of many component factors, or *component causes*. Each component cause is an event or condition that plays a necessary role in the occurrence of some cases of a given disease. For example, the disease might be cancer of the lung and, in the first mechanism in Figure 2–1, factor C might be cigarette smoking. The other factors would include genetic traits or other environmental exposures that play a causal role in cancer of the lung. Some component causes would presumably act in many different causal mechanisms.

Implications of the Causal Pie Model

Multicausality

The model of causation showed in Figure 2–1 illuminates several important principles regarding causes. Perhaps the most important of these principles is self-evident from the model: every causal mechanism in-




Figure 2–1. Three sufficient causes of a disease.