

A Practical Guide to
**CRITICAL
THINKING**



Deciding What to Do AND Believe

DAVID A. HUNTER

 WILEY

A PRACTICAL GUIDE TO CRITICAL THINKING

Deciding What to Do and Believe

DAVID A. HUNTER

Ryerson University
Department of Philosophy
Toronto, Ontario, Canada

 **WILEY**

A JOHN WILEY & SONS, INC., PUBLICATION

This page intentionally left blank

**A PRACTICAL GUIDE
TO CRITICAL
THINKING**

This page intentionally left blank

A PRACTICAL GUIDE TO CRITICAL THINKING

Deciding What to Do and Believe

DAVID A. HUNTER

Ryerson University
Department of Philosophy
Toronto, Ontario, Canada

 **WILEY**

A JOHN WILEY & SONS, INC., PUBLICATION

Copyright © 2009 by John Wiley & Sons, Inc. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.
Published simultaneously in Canada.

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, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permission Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permission>.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Hunter, David A., 1965–

A practical guide to critical thinking : deciding what to do and believe / David A. Hunter.
p. cm.

Includes index.

ISBN 978-0-470-16757-1 (pbk.)

1. Critical thinking. I. Title.

BF441.H86 2009

121—dc22

2008038575

Printed in the United States of America.

10 9 8 7 6 5 4 3

*For Jane,
my beloved and my friend.*

This page intentionally left blank

CONTENTS

Preface	xiii
Note to Instructors	xv
1 The Nature and Value of Critical Thinking	1
1.1 The Nature of Critical Thinking, 2	
Exercise 1.1, 6	
1.2 Critical Thinking and Knowledge, 6	
Exercise 1.2, 7	
1.2.1 Truth, 7	
1.2.1.1 Realism, Relativism, and Nihilism, 8	
1.2.1.2 Relativism and the Argument from Disagreement, 10	
1.2.2 Belief, 13	
1.2.3 Justification, 15	
1.2.3.1 Emotional and Pragmatic Reasons, 16	
1.2.3.2 Epistemic Reasons, 17	
1.2.4 Good Reasons Are Sufficient and Acceptable, 21	
1.2.4.1 When Evidence Conflicts, 22	
Exercise 1.2, 24	
1.3 Critical Thinking and Personal Autonomy, 25	
1.3.1 Belief and Prejudice, 25	
1.3.2 Making Up Your Own Mind, 26	
Exercise 1.3, 29	

- 1.4 Mistakes to Avoid, 31
- 1.5 Practical Strategies, 34
- 1.6 From Theory to Practice: Applying What We Have Learned, 35

2 Clarifying Meaning

39

- 2.1 The Place of Definitions in Critical Thinking, 39
- 2.2 Assertion, 41
 - 2.2.1 Propositions, 42
 - Exercise 2.2.1, 44
 - 2.2.2 Assertion Test, 44
 - Exercise 2.2.2, 48
- 2.3 Constructing and Evaluating Definitions, 49
 - 2.3.1 Slogan, 51
 - 2.3.1.1 Necessary and Sufficient Conditions, 51
 - 2.3.1.2 Necessary and Sufficient Conditions Are Related in Complex Ways, 51
 - 2.3.1.3 Narrow and Broad Definitions, 52
 - 2.3.1.4 Definitions in Practical Life, 53
 - Exercise 2.3.1, 55
 - 2.3.2 Expand, 55
 - 2.3.3 Example, 56
 - 2.3.4 Contrast, 57
 - Exercise 2.3, 58
- 2.4 Thinking Critically about Frameworks, 59
 - Exercise 2.4, 61
- 2.5 Clarifying Beliefs and Problems, 61
 - Exercise 2.5, 65
- 2.6 Technical Definitions, 66
- 2.7 Meaning in Advertisements, 68
 - Exercise 2.7, 70
- 2.8 Mistakes to Avoid, 71
- 2.9 Practical Strategies, 72
- 2.10 From Theory to Practice: Applying What We Have Learned, 73

3 Sufficient Reasons

75

- 3.1 Critical Thinking and Arguments, 75
 - Exercise 3.1, 80

- 3.2 Identifying Premises and Conclusions, 80
 - Exercise 3.2, 85
- 3.3 Dependent and Independent Reasons, 87
 - 3.3.1 The Words Test, 87
 - 3.3.2 The False Premise Test, 88
 - 3.3.3 Circumstantial Reasons, 91
 - Exercise 3.3, 93
- 3.4 Subarguments, 94
 - Exercise 3.4, 96
- 3.5 Evaluating Logical Support, 97
 - Exercise 3.5, 100
- 3.6 Missing Premises, 101
 - Exercise 3.6, 103
- 3.7 Practical Strategies, 104
- 3.8 From Theory to Practice: Applying What We Have Learned, 105

4 Acceptable Reasons

107

- 4.1 Reliable Evidence, 109
 - 4.1.1 Reliability, 109
 - 4.1.2 Undermining and Overriding Evidence, 111
 - Exercise 4.1, 114
- 4.2 Observation, 114
- 4.3 Memory, 117
 - Exercise 4.3, 119
- 4.4 Testimony, 119
 - 4.4.1 Appropriate Testimony, 120
 - 4.4.2 Competent Testimony, 121
 - 4.4.3 Unbiased Testimony, 122
 - 4.4.4 Advertising, 124
 - 4.4.5 News Reports, 126
 - Exercise 4.4, 129
- 4.5 Measurement, 131
 - 4.5.1 Measurement Consistency, 134
 - 4.5.2 Measurement Precision, 135
 - 4.5.3 Surveys, 136
 - Exercise 4.5, 141
- 4.6 Mistakes to Avoid, 143
- 4.7 Practical Strategies, 144
- 4.8 From Theory to Practice: Applying What We Have Learned, 144

5 Reasoning about Alternatives and Necessary and Sufficient Conditions

147

- 5.1 Reasoning about Alternatives, 149
 - 5.1.1 The Meaning of Disjunctions, 149
 - 5.1.2 Denying a Disjunct, 150
 - 5.1.3 False Disjunctions, 151
 - 5.1.4 When Are Disjunctions Acceptable?, 152
 - 5.1.5 Exclusive Disjunctions, 154
 - 5.1.6 Criticizing Reasoning about Alternatives, 156Exercise 5.1, 157
 - 5.2 Reasoning about Necessary and Sufficient Conditions, 159
 - 5.2.1 The Meaning of Conditionals, 160
 - 5.2.1.1 Sufficient Conditions, 160
 - 5.2.1.2 Necessary Conditions, 161
 - 5.2.1.3 Necessary and Sufficient Conditions, 161Exercise 5.2.1, 163
 - 5.2.2 Valid Forms of Reasoning about Necessary and Sufficient Conditions, 165
 - 5.2.3 Invalid Forms of Reasoning about Necessary and Sufficient Conditions, 169
 - 5.2.4 Making It Explicit, 171Exercise 5.2.4, 172
 - 5.2.5 When Are Claims about Necessary and Sufficient Conditions Acceptable?, 173
- 5.3 Reasoning with Definitions and Standards, 174
- 5.4 Reasoning about Causal Conditions, 179
 - 5.4.1 The Meaning of Causal Claims, 180
 - 5.4.2 Reasoning with Causal Claims, 181
 - 5.4.3 When Are Causal Claims Acceptable?, 183
 - 5.4.3.1 Discovering Necessary Causal Conditions, 184
 - 5.4.3.2 Discovering Sufficient Causal Conditions, 185
 - 5.4.3.3 Discovering Necessary and Sufficient Causal Conditions, 186
 - 5.3.4.4 Concomitant Variation, 186
 - 5.3.4.5 Experimenting and Simulating, 187Exercise 5.4, 190
- 5.5 Mistakes to Avoid, 192
- 5.6 Practical Strategies, 194

6 Reasoning by Analogy

195

- 6.1 Reasoning by Analogy, 196
 - 6.1.1 Examples, 196

6.1.2	Is Reasoning by Analogy Valid?, 197	
6.1.3	Relevant Similarity, 198	
6.1.4	When Is an Analogical Claim True or Acceptable?, 200	
	Exercise 6.1, 204	
6.2	Reasoning Using Representational Analogy, 205	
6.2.1	Reasoning with Samples, 207	
6.2.2	When Are Sample Groups Representative?, 208	
6.2.2.1	Sample Size, 208	
6.2.2.2	Random Samples, 209	
6.2.2.3	Self-Selected Samples, 210	
6.2.3	Reasoning with Models and Maps, 213	
	Exercise 6.2, 215	
6.3	Mistakes to Avoid, 217	
6.4	From Theory to Practice: Applying What We Have Learned, 218	
7	Critical Thinking in Action	219
7.1	Thinking Critically about a Discipline, 220	
7.1.1	Identifying a Discipline's Key Concepts, 220	
7.1.2	Clarifying a Discipline's Key Concepts, 222	
	Exercise 7.1.2, 223	
7.1.3	Identifying a Discipline's Sources of Evidence, 223	
	Exercise 7.1.3, 225	
7.1.4	Identifying a Discipline's Modes of Reasoning, 225	
7.2	Critical Thinking Questions, 227	
7.3	Thinking Critically in Your Own Decision Making, 228	
7.3.1	Clarify Your Views, 229	
7.3.2	Clarify Your Reasons, 229	
7.3.3	Show That Your Reasons Are Acceptable and Sufficient, 230	
7.3.4	Identify and Respond to Alternatives, 231	
7.4	Thinking Critically in Discussion, 232	
7.4.1	Ask Open-Ended Clarification Questions, 232	
7.4.2	Withhold Disagreement and Agreement, 233	
7.4.3	Keep Emotional Distance, 234	
7.5	From Theory to Practice: Applying What We Have Learned, 235	
	Appendix A Mistakes to Avoid	237
	Appendix B Practical Strategies	245
	Index	251

This page intentionally left blank

PREFACE

This book has been a long time in the making, and has benefited from the influence of a huge number of colleagues, friends, and family. Here is an inevitably incomplete accounting of some of these debts.

I first began to think systematically about the nature, value, and pedagogy of critical thinking as an assistant professor of philosophy at Buffalo State College and it would be difficult to overstate the influence of my colleague George Hole on my thinking. He is one of the most gifted philosophy teachers I have ever known and I learned a good deal from him on how to teach philosophy. But even more than this, I am indebted to him for the way he so easily mixes philosophy, wit, and good humor in equal parts. I learned more from him than from anyone about how to teach critical thinking, and about the central role it ought to play in education and in a full life. I also owe a great deal to Gerry Nosich, whose work on critical thinking is without equal. Gerry joined us at Buffalo State College as we were designing and implementing a required first-year critical thinking course, and his gentle and wise advice proved invaluable. While with SUNY, I worked on a statewide committee to design a rubric for the assessment of critical thinking. I learned a lot in this time about the importance of teaching critical thinking across the curriculum, and I am especially indebted to Shir Filler.

Since beginning the writing of this book, I have learned a good deal from my new colleagues at Ryerson University, where the philosophy department teaches several sections of a critical thinking course that is

required by students in several programs. I owe special debts to Andrew Hunter, Klaas Kraay, David Ciavatta, Jim Dianda, and Paul Raymont.

I am indebted to Steve Quigley, my editor at Wiley, for gently persuading me to write the book, to Jackie Palmieri, an editorial assistant at Wiley, for gently persuading me to complete it on time, and to several anonymous referees who provided useful feedback on my initial proposal.

I am enormously indebted to my family. I learned as much about how to think critically from my parents as from anyone. They showed me that critical thinking begins at home, and that is a lesson that Miranda and Emily, my wonderful daughters, now champion with exhausting ingenuity.

My greatest and deepest debts, however, are to Jane, whose love and support has never been conditional on sufficient and acceptable reasons. Or on anything else.

NOTE TO INSTRUCTORS

Teaching students to think critically is more about imparting a set of skills and habits than about teaching bits of theory. In developing this textbook, I tried to incorporate several features that I thought would make teaching critical thinking both easier and more effective.

Most significantly, I steered clear of any formal notation aside from the very simplest. It is not that I doubt the value of learning formal logic. In fact, I think that many students not only can benefit from it but can also thrive by studying it. But in my experience there is so much that most students need to learn before they can see the value of mastering a formal system, and so much more benefit they can derive from a non-formal approach to critical thinking. Instead, I tried to think of the text as like an introduction to practical epistemology: offering systematic advice, and lots of practice, on the best way to go about deciding what to believe and what to do.

It is worth noting here that I treat what is sometimes called enumerative induction as a form of reasoning by analogy. It seems to me that using samples to draw a conclusion about an entire group or population just is reasoning by analogy, and that it can be usefully taught as such. I also say, and this perhaps is more controversial, that reasoning by analogy can be valid. Of course, I do not mean that it is formally valid in the way that *modus ponens* is formally valid. Reasoning is valid when it is not possible for the premises to be true and the conclusion to be false. The fact that some reasoning can be known to be valid just from its form alone is, of course, important, and I discuss some of these forms

in chapters 5 and 6. But it is important to keep in mind that not all valid arguments are formally valid (e.g., The table is blue, therefore it is colored), and not all arguments that are formally invalid are really invalid (e.g., If Jones is a male, then he is a bachelor; Jones is a bachelor; so, Jones is a male.) Judgment is always needed, it seems to me, in assessing the strength of a piece of reasoning, and this judgment is better taught by focusing on the idea of validity itself. I also think that what I say in Chapter 6 makes a reasonable and pedagogically responsible case for my view that reasoning by analogy can be valid.

I had originally planned to dedicate a chapter to thinking critically about what to do. But I worried that much of it would simply repeat points that had been made earlier and, in so doing, would make deciding what to do seem like a lesser cousin to deciding what to believe. As I worked (and then re-worked) the first six chapters, it seemed to me that I could elegantly discuss deciding what to do as we went along, when the topic at hand seemed relevant. I have thus included several “boxes” discussing various aspects of deciding what to do.

The book includes several other kinds of boxes as well. Some identify important mistakes that a good critical thinker ought to avoid. Some provide summaries of the discussion in the body of the text. Some offer examples of critical thinking across the curriculum. Some offer practical tips and rules of thumb. All are intended to make the text more readable and the concepts and skills more accessible.

I also decided that rather than dedicate a chapter to informal fallacies I would discuss them in what struck me as their proper context. It seems to me that there is no easy way to organize the different kinds of mistakes into a small number of categories without distorting their differences or exaggerating their similarities. Some of the mistakes have to do with clarifying meaning; others with ascribing views to others; some with assessing evidence; others with assessing validity. Several mistakes can occur at several otherwise quite distinct stages in deciding what to do or to believe. Rather than try to force the various mistakes into artificial categories, it seemed to me better to discuss them as we went along. For easy reference, though, I have collected them all in an appendix at the end of the book.

Careful training and repeated practice are crucial to learning any skill, and critical thinking is no exception. I have tried to include a large and varied collection of exercises. But I strongly encourage you to bring your own exercises to class and to encourage your students to seek out arguments and reasoning to share during the class time. In my experience, students learn far more when they are required during class time to participate in the construction, analysis, and assessment of

examples of reasoning about what to do or believe. I have included, at the end of most of the chapters, exercises that are specially designed to help students transfer the concepts and skills they are learning to other corners of their lives. My thought is simply that there is little point in teaching someone to think critically if they see no place for it at home, in their own discipline, or at work. Over the years I have experimented with all of these exercises, making adjustments as I went along. The exercises are in a form that I find to be both effective and not overly intrusive. But I encourage you to adjust, alter, add, subtract, and modify as you see fit. The important thing is to find ways to help students see that they are learning skills and concepts that have application and value after the final exam.

This page intentionally left blank

1

THE NATURE AND VALUE OF CRITICAL THINKING

This book is a practical guide to critical thinking. It might seem unnecessary to be reading a guide to something you do all the time and are probably already pretty good at. When I tell people that I am writing a book on critical thinking they sometimes tell me that they consider themselves to be very good critical thinkers. At the very least, they say that they consider critical thinking to be very important. I am sure that they are right on both counts. We think critically a good bit of the time, and on the whole we do it pretty well. Still, I think there is always something to learn from thinking hard about what one is already good at.

In this chapter, we will explore the nature and value of critical thinking. We will ask what critical thinking is and how it differs from other kinds of thinking. We will explore what it *means* to think critically; what makes that kind of thinking *critical*. As part of this, we will consider whether critical thinking varies from one discipline to the next. Is critical thinking in geology different from critical thinking in design or the humanities? We will see that while the concepts, methods, and standards may differ from one discipline to the next, there is a basic essence or core of critical thinking that remains the same across all disciplines. Whether one is doing chemistry, design, astrology, or philosophy, there

are common standards that you should strive to maintain, and practical strategies to help you make sure that you do. This book is designed to introduce you to this essential core of critical thinking while at the same time providing you with the tools you need to identify the concepts, methods, and standards distinctive of different disciplines.

Once we have said what we mean by critical thinking, we can then ask what place this kind of thinking does or should occupy in our daily lives, both in and out of the classroom. When is it appropriate to think critically, and are there some parts of our lives where critical thinking tends to dominate or where it tends to be ignored? We will see that critical thinking is appropriate whenever we are trying to decide what we ought to believe about some matter of fact or whenever we are trying to decide what to do or what course of action to adopt. In short, critical thinking is needed whenever we reason about what to believe or what to do.

Finally, and perhaps most importantly, we will ask why being a critical thinker *matters*. What makes critical thinking valuable? Why should we engage in it? We will see that being a critical thinker is valuable for several reasons. Perhaps most obviously, thinking critically about a question or problem can help one get the right answer or solution. By thinking critically about what to believe or what to do we increase our chances that our beliefs will be true and our actions effective. Thinking critically may not guarantee that you get the right answer; however, a good case can be made that unless you think critically you will get the right answer only by luck, and relying on luck is not a wise policy. But critical thinking has a deeper value than just its ties to truth. Critical thinking is also closely tied to one variety of freedom. By thinking critically, one can make up one's own mind and making up one's own mind is essential if we are to be the master of our own lives. Critical thinking, we will see, is essential to personal autonomy.

1.1 THE NATURE OF CRITICAL THINKING

There are many definitions of critical thinking, but Robert Ennis, one of the leading researchers on critical thinking, offered the following definition many years ago and it remains, to my mind, the best of the bunch: "Critical Thinking is reasonable, reflective thinking that is aimed at deciding what to believe or what to do."¹

¹Ennis, R. H. "A Taxonomy of Critical Thinking Skills and Dispositions," in *Teaching Thinking Skills: Theory and Practice*, ed. Joan Boyloff Baron and Robert J. Sternberg (New York: Freeman, 1987), pp. 9--26.

We can see that there are several elements to this definition, so let us look at them one at a time, starting with the last one.

Critical thinking is thinking that is *aimed at deciding what to believe or what to do*. Deciding *what to believe* is a matter of deciding what the facts are, figuring out what the world is like, or at least what some little corner of it is like. We make these kinds of decisions when we decide whether it is raining out or sunny, whether the Blue Jays stand a chance this year, whether the kids will put up with another meal of macaroni and cheese, whether the movie was as good as its billing, whether the restaurant has gotten better over the years, or whether we should trust what our teachers tell us. In deciding what to believe on some matter we take a stand on it. If it is a decision on a factual matter, like the decision about the weather or about the Blue jays, then we take a stand on what the facts are. If it is a decision on an evaluative matter, like the one about the movie or the restaurant, then in deciding what to believe we are taking a stand on what is good or better. In either kind of case, critical thinking is aimed at helping us to make those kinds of decisions about what to believe.

Critical thinking is also *aimed at decisions about what to do*. Deciding what to do really has two parts. First, one has to decide what to value or to strive for. This is a matter of deciding on one's goals or end. Then, one has to decide how best to achieve that end. This is a matter of deciding on the best *means* to that end. Should I go for a run now or keep working on my book? Should I spend my savings on a new car or continue using my beat-up one? Should the city spend its limited resources on building a new bridge? Should the country move towards a universal health care plan? Should I tell the truth when my friend asks me about her boyfriend? Should I give to charities? Usually we decide what to do on the basis of what we already value or on what we already think makes for a good life. I decide to go for a run instead of continuing to work on this book because I feel that running and staying in shape is an important part of my life. I decide to tell the truth to my friend about her new boyfriend because I value honesty in my friends and want them to consider me trustworthy. But sometimes, deciding what to strive for or what goals to pursue requires first deciding what one will value, what kind of person one wants to be, what kind of life one wants to lead. In deciding whether to pursue graduate school in philosophy, I had to make a decision about to value, about what kind of shape I wanted my life to take. Decisions about what to value are among the most difficult and profound decisions we can make. Critical thinking can help us to make these kinds of decisions. But once we make them, once we decide what we want our life to be like, we still

need to decide what the best way is to make our life that way. Once we choose the ends, we still need to decide on the means. Here too, critical thinking can help.

According to Robert Ennis' definition, critical thinking is **reasonable** thinking. This is so in several respects. First, critical thinking is reasonable thinking because it is sensitive to methods and standards. If we are trying to decide what to make for dinner or whether the Blue Jays stand a chance this year, there are various methods we should use and standards we should keep in mind. If we try to make up our minds on these topics without relying on those methods or obeying these standards we will fail to be thinking critically about the topic. Part of what makes critical thinking *critical* is that it is governed by rules and methods. This does not mean that there is not plenty of room in critical thinking for judgment and flexibility. In fact, as we will see in a moment, part of what makes critical thinking different from other kinds of thinking, such as arithmetical calculation, is that there is room for judgment and a case-by-case flexibility. Still, it is essential to critical thinking that in thinking critically about what to believe or do we rely on methods and are subject to standards. We will spend lots of time in the following chapters learning about what these methods and standards are.

Critical thinking is *reasonable* in another and deeper sense. Critical thinking about what to believe or what to do is reasonable in that it demands that we have reasons, and preferably good ones, for making the decisions we do. The aim of critical thinking is not simply to make a decision on what the facts are or what to strive for. In a way, it is easy to make such decisions. What is hard is having good reasons for the decisions we make. It is not enough to decide to believe that it is sunny out; one has to have good reason to decide this. Likewise, it is not enough just to decide to value honesty or justice; one has to have good reason for this decision. So critical thinking is reasonable in that it demands that we have reasons, and preferably good ones, for making the decisions we do. We will be spending a lot of time in what follows exploring what makes something a *good reason* to believe or to do something.

Finally, Ennis says that critical thinking is **reflective**. We can see what he has in mind if we contrast critical thinking with arithmetical calculation. There is no doubt that calculating the square root of a large number is a kind of thinking and no doubt that it is thinking that is sensitive to methods and standards. In this respect, arithmetical calculation is like critical thinking. But when one calculates a number's

square root, one does not need to think about the methods one is using. One simply uses the formula to get the right answer. In this kind of case, the problem at hand (finding the number's square root) is pretty straightforward: it is perfectly clear from the beginning what is to count as the right answer and what the best means is of finding it. The same is true for many kinds of decisions we make in our daily lives. But some problems are **open-ended**. A problem is open-ended when it is not clear from the outset what would count as a solution to it. In such cases, progress may require thinking hard about the problem itself, and not just calculating an answer to it. To solve it, we may need to analyze the problem into parts, and we may need to think about the best method to use to find a solution, and while we employ that method we may need to be thinking about whether we are employing it correctly. We may even need to adjust the method or even develop one from scratch. I'll have more to say later about open-ended problems and no doubt the line between straightforward ones and open-ended ones is not hard and sharp. Calculating a square root the first few times requires a good deal of reflection even when one does have the formula; and deciding whether it is raining or sunny is usually as straightforward as looking out the window. Still, the contrast should be clear. Critical thinking is *reflective* in the sense that it involves thinking about a problem at several different levels or from several different angles all at once, including thinking about what the right method is for answering or solving the problem.

One of the chief virtues of this definition is that it does not restrict critical thinking to the study of **arguments**. An argument is a series of statements some of which (the premises) are meant to provide logical support for another (the conclusion). Because we can and often do formulate our reasons for believing or doing something in the form of an argument, critical thinking is surely concerned with arguments. In later chapters we will discuss some strategies and standards for analyzing and evaluating arguments. But the notion of an argument does not always fit naturally across the curriculum. It is hard to see how reasoning about experimental design or about statistical sampling fits the paradigm of an argument. What is more, evaluating reasons for believing something involves assessing their acceptability and their meaning, and neither of these tasks is ordinarily considered argumentation. It is, of course, possible to stretch the ordinary concept of an argument or of argument analysis to include all these different aspects of critical thinking. But this definition captures them all without artificially extending our ordinary words.

EXERCISE 1.1

1. Short-answer questions:
 - a. In what sense is critical thinking reflective?
 - b. What makes critical thinking reasonable thinking?
 - c. Why is arithmetical calculation not a kind of critical thinking?
 - d. Does critical thinking have to be “critical” in the sense of being negative or skeptical? Explain, using an example.
2. Which of the following activities involves critical thinking? If an activity does not involve critical thinking, identify which element in critical thinking is missing.
 - a. Riding a bike
 - b. Watching the news on TV
 - c. Doing laundry
 - d. Ordering coffee at a local coffee shop
 - e. Planning a vacation
3. Identify five activities you do on a daily basis that do not involve critical thinking. Identify two or three activities that you do on a daily basis that would be improved by thinking critically about them, and explain how thinking critically would improve it.
4. Now that you know what critical thinking is, list five reasons why it is good to think critically.
5. List five possible obstacles to thinking critically. Describe one strategy for overcoming each obstacle.

1.2 CRITICAL THINKING AND KNOWLEDGE

We have been discussing what critical thinking is and we can now explore why it matters. As I said at the outset, critical thinking is valuable for two main reasons. First, thinking critically increases our chances of gaining **knowledge**, and knowledge is valuable. Second, thinking critically is essential to making up one’s own mind about what to believe or what to do, which is essential to being **autonomous**, and being autonomous is valuable. We will discuss knowledge in this section, and autonomy in the next.

We have seen that critical thinking is thinking that is aimed at deciding what to believe or to do. But ideally we want more than just to have

an opinion about the facts; we want to *know* what they are. When we check the weather, our goal is not just to reach a decision about whether it is sunny or not; we want to come to *know* whether it is sunny or not. We want to *know* whether the city ought to spend its scarce resources on building a new bridge. We want to *know* whether HIV causes AIDS all by itself or only in conjunction with other factors. So critical thinking is really aimed at knowledge. But what is knowledge? What is it to know something? By answering these questions we can get quite a bit clearer on what critical thinking is and why it is valuable.

EXERCISE 1.2

We can start with an exercise. Make a chart with three columns. In the first column, list things that we, either individually or as humans in general, know for a fact. In the second column, list things that we can know, but currently do not know. In the third column, list things that we do not and probably cannot ever know about. These can be particular facts or kinds of things. The more variety you can provide in each list the better. (Include something in one of the columns only if you are fairly sure that everyone else in your class would also include it in that column. This will avoid controversy from the start.) When you have the Knowledge Chart completed, compare the items in the first and second column and try to identify the relevant differences? What is lacking in the items in the second column that prevents their being in the first column?

The traditional definition of knowledge developed by philosophers says that knowledge is justified, true belief. According to this definition there are three elements to knowledge. We can look at each in turn. Then we will ask how the three elements are related to one another. Let's start with truth.

1.2.1 Truth

It would be ideal at this point in our discussion to provide a clear and precise definition of truth. I do not mean just a listing of all the truths that there are, though such a list would be valuable. We already know some of what such a list would include. It would have to include the truths that Barack Obama is the 44th President of the U.S., that a virus causes the flu, and that the Earth orbits the Sun. And we know what things we should leave off that list: it is not true that fish are birds, it is not true that $2 + 2 = 27$, and it is not true that George Washington

was president of France. It would probably be impossible, or at least really hard, to make a complete list of all the truths. But even if we could, making such a list would not be the same as giving a definition of truth. To give a definition of truth we would have to say what it is for something to be true. We would need to say, in a general sort of way that would apply to every case, what *makes* something true. I do not have any idea how to do this. Nor, I think, does anyone else. Or rather, the only definition that I know of is not very helpful: a statement is true just in case it corresponds with the facts. This is not that helpful because the notion of corresponding with the facts is not clearer than the notion of truth itself. Thankfully, though, we do not really need a definition of truth. For our purposes it will be enough to contrast three attitudes we might take to some subject matter: **realism, relativism and nihilism.**

1.2.1.1 Realism, Relativism, and Nihilism A **realist** about some subject matter is one who thinks (i) that there are truths in that area and (ii) that what those truths are is independent of what anybody thinks they are. In saying that those truths are independent of what anybody thinks that they are, I mean that they would be true even if nobody knew or even believed that they were true. The truth, as it were, is simply “out there.” Because she thinks that truth is independent of our beliefs, a realist thinks that it is possible (even if it is highly unlikely) that we could all be totally mistaken about or ignorant of the facts in that subject area. She might even think that the facts are beyond our understanding, that no matter how hard we tried or for how long, we simply cannot come to know those facts. Of course, being a realist does not mean that one has to be skeptical or doubtful about whether we do know anything about that subject matter. One can be a realist about a subject matter and still be quite confident that we know a lot about it. Being a realist simply requires thinking that the facts in that subject area are not determined by or dependent on our beliefs about them. They are what they are, regardless of what we might think that they are.

A **relativist** about some subject matter holds that (i) there are truths about that area but (ii) that what they are depends (in some way or other) on what we (or someone) take those truths to be. The relativist and realist agree that there are truths or facts of the matter in that area, but they differ over how those truths or facts are related to our beliefs about them. The relativist insists that those facts are what they are because of our beliefs about them, whereas the realist insists that our beliefs have no bearing at all on the facts themselves. The relativist

maintains that had our beliefs or our natures been different, then the facts might have been different too. The facts somehow depend on us. This means that on a relativist's view of some subject matter, it is in a certain way impossible for us to be wrong or ignorant of the facts in that area, since our beliefs about what the facts are is at least part of what makes them the facts. We cannot go too far wrong in trying to know that subject matter because we play an essential role in making the subject what it is.

There are different versions of relativism, differing in terms of whose beliefs play the role of determining what the facts are. A **subjective relativist** about some topic is one who thinks that the facts in that area are whatever any one individual takes them to be. She might express this idea by saying things like: "Well, that might be true for you, but it is not true for me." A **social relativist**, by contrast, holds that the facts in that area are whatever the majority (or some weighted majority) of the society or culture takes them to be. "It is true for us, even if it is not true for you or for them." What is common to all versions of relativism is the idea that the facts are in some way dependent on our beliefs about them; that, in one way or another, the facts are what they are because we are the way we are.

A **nihilist** about some subject matter holds that there are no truths at all about that subject matter. There are, on the nihilist's view, no facts to be right or wrong about. It is not that the facts depend on us in some way; there are no facts at all (aside from the fact that there are no facts). There is no such thing as truth in that area. Since there is no such thing as truth in that area, there is also no such thing as knowing the truth, and not because we are incapable of coming to know it, but because there is nothing there to be known at all. The nihilist thus disagrees with both the realist and the relativist, though as we just saw, the realist and the relativist also disagree with each other.

One could be a realist about one subject matter and a relativist about another and a nihilist about a third. One might think, for instance, that realism is the proper attitude to take to particle physics or to human history, but think that nihilism is the right attitude to take towards the nature of Santa Claus. Or one might be a realist about human biology but a relativist about humor, thinking that while the facts about our biological natures are independent of our beliefs about them, whether something is funny or humorous does depend on whether we find it or believe it to be funny. Indeed, it is hard to see how we could possibly all be wrong about whether some joke is amusing. Maybe what makes something funny is simply that we all (in normal conditions) believe it to be funny. If so, then perhaps relativism is right about humor.

One cannot take two or all three of those attitudes to one and the same subject. One could not be both a realist and a nihilist about, say, particle physics. For this would mean holding (as a realist) that there are facts about particle physics while also (as a nihilist) denying that there are facts about particle physics. But this is incoherent. Realism and nihilism about some subject matter are contraries of one another: they cannot both be true, though they could both be false. Likewise, one could not be a relativist and a realist, or a relativist and a nihilist about one and the same subject matter. But in principle one could, and I think we in fact sometimes do, take different attitudes to different subject matters or topics.

One has to have good reasons for being a realist, relativist, or nihilist about some subject matter. It is not enough simply to decide or declare that one will be a relativist about, say, particle physics or geometry, or a nihilist about morality and geography, or a realist about humor and beauty. One has to be able to provide good reasons for thinking that one is taking the appropriate attitude to that subject. If one is a realist about particle physics but a relativist about humor, then one has to be able to explain what the difference is between those subject matters or about our relations to them that warrants taking those different attitudes to them. The explanation cannot simply be that the facts about particle physics are independent of us whereas those about humor are not. To say this is simply to express your attitudes, not to justify or explain them.

1.2.1.2 *Relativism and the Argument from Disagreement* Relativists about some subject matter sometimes try to justify their attitude by pointing to the fact that there is little or no agreement among otherwise well-intentioned and sincere people about what the facts are in that subject matter. Relativists about morality, for instance, point out that there is considerable disagreement among sincere people about just what our moral duties are, or about how to balance competing moral demands. And they suggest that the existence of this kind of disagreement lends support to their relativism. We can formulate this reasoning for relativism about morality as an argument:

- i. There is only considerable sincere disagreement over moral facts.
- ii. If there is only considerable sincere disagreement over the facts in some area, then relativism is true of that subject area.
- iii. So, relativism is true of moral facts.

Let us call this argument for moral relativism, the **Argument from Disagreement**. It would be easy to transform it into an argument for any kind of relativism. We could get an argument for relativism about humor by replacing the word “moral” with the word “humor.” But let’s focus on this argument, since the main lessons will apply across the board.

The Argument from Disagreement has an important logical property. It is **valid**. This means that if the premises (i.e., claims (i) and (ii)) are true, then the conclusion (i.e., claim (iii)) would have to be true too. In other words, it is not possible for those premises to be true and yet for the conclusion to be false. If the premises are true, then they constitute a conclusive proof that moral relativism is true. We will have much more to say about validity in Chapter 3. But for now, it is enough to note that when an argument is valid, the only question that needs to be considered in evaluating it is whether the premises are true. So let’s consider each premise.

The first step in deciding whether a premise is true is to make sure that we know exactly what it means. This is a bit difficult in the case of the *Argument from Disagreement’s* first premise because it is not very clear what “considerable” means. How much disagreement counts as “considerable?” Does everyone’s opinion count equally in deciding when moral disagreement is considerable, or are there moral experts whose opinions matter more? What if the moral theorists all agreed but that everyone else held different opinions? Would premise (i) be true in that case? These are difficult questions about just what claim premise (i) is making, and it is not clear how best to answer them. But let us set aside these questions for now. Let us suppose that we had some good method for measuring when disagreement is considerable. There is another aspect of the meaning of the first premise that we need to pay close attention to. It says more than just that there *is* considerable moral disagreement; it says that there is *only* considerable moral disagreement. That means that there is not also considerable sincere *agreement* over the moral facts. This will be important.

Now that we are pretty clear about what that first premise means, let’s see whether we have good reason to accept it. Is it *true* that there is only considerable sincere disagreement over moral facts? It certainly does seem to be true that there is disagreement over moral facts. Different societies have held different views about what morality requires or permits. There are sometimes disagreements among people in our own country or even within our own family about morality. So it is hard to deny that there is a disagreement over morality. But many

researchers have pointed out that even though different societies disagree about some moral claims, there is also often quite broad and deep agreement about others. For instance, even though different societies have different views about which marital and sexual practices are morally acceptable, every society thinks that sexually assaulting one's own children for pleasure is morally wrong. And even though we might disagree with our friends over whether it is morally wrong to be drunk, we probably all agree that it is morally wrong to drive drunk. So it is not obvious that there is *only* considerable sincere moral disagreement; there also seems to be considerable sincere moral agreement. Indeed, it might even be that while there is a lot of disagreement about just what it is that morality requires, there is at the same time just as much or even more agreement about what morality requires. This shows that it is not clear that the first premise in the Argument from Disagreement really is true.

What about the second premise in the Argument from Disagreement? It says that if there is only considerable sincere disagreement over the facts in some area, then relativism is true of that area. Is this true? We can begin by noting that the existence of disagreement would not all by itself show that relativism is true of an area. There is lots of disagreement among physicists over the fundamental features of our universe. But this does not incline us to be relativists about physics. Indeed, this amount of disagreement is exactly what we expect from a subject as complex and difficult to understand as physics. One reason we continue to be realists about physics is that there is also considerable agreement (at least among experts) about the physical facts, in fact there is far more agreement than there is disagreement. Moreover, as hard and complex as physics is, it still seems that we are making progress. But what if after a long and exhaustive attempt to reach agreement in some field, we found only widespread and sincere disagreement with little or no agreement at all and no sense that progress was being made? (This is not, as we have seen, the situation with respect to morality, since there is considerable agreement about moral facts, even though there is also considerable disagreement. Question: Is there also reason to think we are making progress in morality?) Would this justify being a relativist about that subject matter? Or would it instead justify being a nihilist about that subject matter? If we could never reach any substantial level of agreement, should we say that the facts depend on us, or should we say that there are no facts? Under what conditions would it be right to conclude, with the nihilist, that there are no facts at all, that we have been misled somehow into thinking there are facts when there really are not? I am not sure

how to decide this question. I find it hard to know when to be a relativist instead of a nihilist. In any event, it seems clear that the existence of nothing but considerable sincere disagreement in some subject matter would not necessarily show that relativism is true of that area. So it is not obvious that premise (ii) in the Argument from Disagreement is true.

We have seen that there is good reason to doubt the truth of both premises in the Argument from Disagreement. It is not true that there is only considerable disagreement about moral facts. And even if there were, it is not clear that this would show that relativism is true of morality. So the Argument from Disagreement does not show that moral relativism is true; the argument is not successful. But the fact that the argument is not successful does not show that moral relativism is false. The conclusion of a bad argument might still be true. All we have shown is that one set of reasons for believing in moral relativism are not good ones. It might be that there are other, much better reasons for thinking that moral relativism is true. And of course it might be true even if we cannot find any reasons to believe that it is true. Still, as a good critical thinker we ought not to believe that moral relativism is true unless we have good reasons to believe that it is true. The same is true, of course, for the realist or the nihilist; we all need to have good reasons for our beliefs.

Nonetheless, the realist might have a slight *methodological* advantage over both relativism and nihilism. It is sometimes suggested that relativism and nihilism are obstacles or impediments to critical thinking. I do not think this is true. What is true is that *unjustified* relativism and *unjustified* nihilism are impediments to critical thinking. One should not be a relativist or a nihilist without good reasons. But perhaps in the absence of convincing reasons to be a relativist or a nihilist, we ought to work under the assumption that realism is the proper attitude to take. Maybe realism is the proper *default* view to take, so long as we take it with an open mind, until we are shown that it is wrong. Perhaps it is better to err on the side of realism than to err on the side of nihilism or relativism. In any event, the critical thinking strategies and standards we will be discussing in the following chapters will assume that realism is the appropriate attitude to take. We will assume that truth is independent of our beliefs.

1.2.2 Belief

The traditional philosophical analysis of knowledge says that knowledge is or requires justified true belief. This means that to know

something you also have to believe it. Sometimes we contrast what we know with what we merely believe to be the case, and sometimes when we talk about our beliefs we have in mind our views on moral or religious topics, where it is hard to find general agreement. If you were asked to list your beliefs, you might describe your views on God, happiness, justice, but not include your views on the day's weather, on your favorite sporting team's recent performance, or on arithmetic. It even sounds a bit odd to say that I believe that $2 + 2 = 4$. It is tempting to say, "I don't believe it; I know it." But I think that we find this odd to say because it leaves the mistaken impression that we do not also feel quite confident that we know it. To say that I *believe* that $2 + 2 = 4$ would be to say something weaker than what I could say, and that is what makes it a misleading way to put it. But it might be true that I believe it, even if it would be misleading to say it. In any event, in this book we will follow the philosophical tradition and assume that to know something you must also believe it. Our real concern is with justification anyway and not with belief. Critical thinking is concerned with the *kinds of reasons* that are needed to know something.

There is another reason to follow the philosophical tradition here. If we separate off too sharply what we know from what we believe, then we run the risk of overlooking the fact that even our religious and moral beliefs need to be based on good reasons. It is true that we have and should cherish **freedom of belief**. Being able to form our beliefs free from outside interference and coercion is fundamental to human fulfillment. We should be permitted to make up our own minds on religious and moral topics. This means that there are limits to the kinds of criticism that can be directed at our beliefs on such topics. But, and this is the crucial point, it does not mean that there are no epistemic standards against which our beliefs on these topics can be assessed. After all, freedom of belief is not restricted to moral and religious topics. We should also be free to make up our own minds about the weather, arithmetic, human evolution, and the best use of scarce public resources. Our beliefs about the weather and about human evolution still need to be based on good reasons, even if we ought to be allowed to make up our own mind on those topics. So freedom of belief does not mean that we don't need to have good reasons for our beliefs. In fact, as we will see in the next section, having good reasons for our beliefs is essential to genuinely making up our own minds. Critical thinking is appropriate not just when we think about the weather or about public policy. The standards and methods that are central to critical thinking are also appropriate when we decide what to believe about God, justice, or morality.

BOX 1.2.1 FREEDOM AND RESPONSIBILITY

We have and value freedom of action, as well as freedom of belief. It is important to us that it be, in some sense, up to us what we do and where we go. But this freedom is *limited*: no matter how much I might want to or how hard I try, I'm not free to jump to the moon or grow ten inches in a day. And freedom of action brings great *responsibility* too: I am not free to torture or hurt people for the fun of it.

Are there also limits to what you can believe? Could you now, at this very instant, voluntarily *make* yourself believe that $2 + 2 = 27$, or that the Earth really is at the center of the solar system? Or are your beliefs not under your immediate voluntary control? Would you like them to be?

Are there also *responsibilities* that come along with having beliefs? Would it be irresponsible for you to believe that the Earth is at the center of the solar system? Why or why not? What would make it irresponsible?

1.2.3 Justification

We have seen that knowledge is or requires a justified true belief. To say that a belief is justified is to say that it is based or grounded in good reasons, that the believer has adequate or satisfactory reason to hold or to sustain her belief. But there are lots of different kinds of reasons to believe something, and it is worth distinguishing some of them so that we can focus on the kinds of reasons that critical thinking is concerned with. It will help to have an example, so let us suppose that Jones believes that humans evolved from other living species, in something like the way current theories of evolution describe. We can ask three questions: (i) What kinds of reasons might Jones have for believing this? (ii) What kind of reasons is critical thinking concerned with? (iii) What is it for reasons of that kind to count as *good* reasons?

We should start by noticing a distinction between **producing** reasons and **sustaining** reasons. The producing reasons are the ones that made Jones believe it in the first place, whereas the sustaining reasons are the ones that his belief is now based on. The producing reasons need not be the same as the sustaining reasons. Perhaps Jones first came to believe that humans evolved from other species because he heard it on a TV show that he has now long forgotten about, but continues to believe it because of the evidence he has since acquired in various

science classes. In that case, the producing reasons are not at all the sustaining reasons. It is of course possible for the producing reasons to also be the sustaining reasons. No doubt for the first few days after watching that show, the reasons that produced his belief also sustained it. But this does not have to be the case. I suspect that for many of our beliefs, the reasons that we had for forming them are not those that now sustain them. There is nothing wrong with this. Indeed, it is to be expected, I think, that as our evidence changes and grows this will affect the reasons we have for what we believe. But it is still important to keep the difference in mind when we are asking why someone believes something, since criticizing the reasons he originally had might be beside the point if those are no longer his reasons.

1.2.3.1 Emotional and Pragmatic Reasons As I said at the outset, there are many different kinds of reasons to believe something. One can have **emotional** reasons to believe something. Maybe Jones believes that humans evolved from other species in part because believing it helps him feel at one with his natural environment, and this feeling brings him a deep sense of connectedness and meaning. Giving up that belief might cause a sharp emotional pain or rupture. Or maybe he believes it because he knows that believing it upsets his religious father, and he derives satisfaction in being rebellious. Or maybe that belief fits into a larger web of beliefs he has about his place in the universe, and giving it up would damage the integrity or coherence of that web of belief in a way that would be hard for him to accept. Some of our beliefs are simply so fundamental that giving them up would cause a huge and unpleasant upheaval in our personal worldview, and the desire to avoid this can itself be a reason to keep the belief. Some theorists have suggested that emotional reasons play a fundamental role in producing or even sustaining our moral or religious beliefs. Perhaps Jones' belief that lying to others is wrong stems from feelings of guilt he has when he lies, or from feelings of shame he has when he has to admit to others that he has lied. Perhaps he believes that God exists partly because it brings him deep comfort. Moral and religious beliefs do not have to be produced or sustained by emotional reasons. And I suspect that emotional reasons play a role in many of our ordinary "factual" beliefs. It is important to us to feel balanced, and sometimes the need to continue to feel balanced plays a role in explaining why we continue to believe what we do.

We have been considering emotional reasons to believe something that involves only the believer himself. But one can also believe something because of the way that belief relates one emotionally to one's

community, culture, or heritage. Having a strong sense of community and tradition is extremely important to us, and we should not underestimate the way it can influence and shape our view of the world. Perhaps Jones identifies with the scientific community and tradition and thinks that not believing in human evolution would force him to break with that community and that this break would be bad or painful. It is certainly true that many of the practices we currently have are sustained, at least in part, in order to strengthen and nourish strong community bonds. Sometimes, our practices and beliefs are so fundamental not only to our own personal worldview but to our cultural and ethnic heritage that it is hard to see them as anything but natural and inevitable. It may seem to us that not maintaining them would be a kind of lunacy. (Sometimes, it is only by studying foreign practices and traditions that one can really appreciate and even identify one's own heritage and practices for what they are.) In this kind of case, it might be impossible to even question the beliefs or practices without causing substantial emotional pain.

We can also have more purely **pragmatic** reasons to believe something. We might believe something because believing it makes it easier for us to achieve our goals or objectives. It might be that abandoning the belief would not cause us serious emotional pain of any kind, but that we find that maintaining the belief simplifies some part of our practical life. It is easier to get along if we believe it than if we question it, and so we continue to believe it.

1.2.3.2 Epistemic Reasons We have been discussing reasons to believe something. But so far we have not discussed reasons to believe that something is true. Let us call reasons of that kind, ones that indicate that what we believe is true, **epistemic** reasons. Emotional reasons and pragmatic reasons are not epistemic ones. Even if it is true that abandoning some belief would cause substantial pain or practical difficulty, it does not follow that these reasons for sustaining the belief are also reasons to think that the belief is true. This is clear, I think, in the case of Jones' belief that humans evolved from other species. The emotional or pragmatic reasons he has to believe that have nothing at all to do with whether the belief is true. Indeed, the truth of his belief has nothing at all to do with his emotions, or his community or even with him. If it is true that humans evolved from other species, then this is true regardless of whether Jones even exists. Whether it is true depends on events that occurred long before he was born. Epistemic reasons are reasons to think that a belief is true or accurate, that it captures the facts properly, and they need have no special bearing on

our emotions or practical challenges. Indeed, as we all know, sometimes the truth is painful or uncomfortable.

Epistemic reasons are at the heart of critical thinking. Think back to the traditional philosophical definition of knowledge as justified true belief. Since knowledge requires *true* belief, the kinds of reasons involved in justification are epistemic ones, not emotional or pragmatic ones. The requirement that to know something one's belief must be justified means that one must have good epistemic reasons for the belief. One must have enough of the right kind of evidence. Basing or sustaining a belief on emotions or on practical considerations cannot lead to knowledge, since these kinds of reasons to believe something are the wrong kind. To know whether humans evolved from other species it is not enough to have strong emotional or pragmatic reasons; one must have strong reasons for thinking *that it is true* that humans evolved from other species. This does not mean that one cannot also have emotional or pragmatic reasons. Jones's belief that humans evolved from other species might be justified enough for knowledge even if it is sustained in part by emotional or pragmatic reasons, so long as he also has sufficient epistemic reasons to believe it. But if one is striving for knowledge, then one cannot rest content merely with emotional or pragmatic reasons, since they have nothing essentially to do with whether the belief is true, and truth is essential to knowledge. A belief that is based solely on emotional or pragmatic reasons cannot possibly count as knowledge, even if the belief is true. Knowledge requires strong epistemic reasons.

BOX 1.2.3A CRITICAL THINKING AND THE PRACTICE OF MEDICINE

In an article in *The New Yorker*, Dr. Jerome Groopman wrote about how doctors sometimes let emotions get in the way of their examinations. He described a case in which he missed a patient's serious infection because he did not want to embarrass his patient by doing a thorough physical examination. Had he looked carefully, he would have found a serious infection. Luckily, another doctor discovered the infection and it was treated. Groopman's mistake, in this case, was not that he based his beliefs on his emotions, but that he allowed his emotions to get in the way of performing the kinds of tests and examinations he knew were needed before deciding whether the patient was healthy. He allowed himself to form a belief that he

knew was based on incomplete evidence. The consequences of this mistake might be just as bad as the consequences of forming beliefs on the basis of emotions. He wrote that this case illustrates an important lesson, neatly summarized by his friend Pat Croskerry: “Currently, in medical training, we fail to recognize the importance of critical thinking and critical reasoning. The implicit assumption is that we know how to think. But we don’t.”

BOX 1.2.3B DECIDING WHAT TO DO: DECIDING ON MEANS AND DECIDING ON ENDS

Deciding what to do involves two separate decisions. The action’s intended goal—its *end*—is one thing and the steps to achieve that goal—the *means*—are quite another. Here are some examples:

I will enroll in university in order to get an education.

I’ll dedicate all of June and July to writing my book in order to get it done.

I’ll put a pot of boiling water on in order to make dinner.

Thinking critically about what to do requires having reasons to pursue those ends and reasons to choose those means.

Reasons for pursuing some end are reasons for thinking that the end is good, or valuable, or worthwhile. I decided to write this textbook because I believed that writing it would be a good thing to do, and I had reasons for this. You decided to go to college or university because you thought it would be a good thing to do, and you surely had some reasons to think that. Reasons to think that something is good are a special kind of reason, and we will look at them in more detail in a later chapter.

Once you decide on your goal or your end, you need to decide how to make it happen. This is deciding on the means to achieve that end or goal. Reasons to adopt some means are reasons for thinking that those means will succeed. I decided to dedicate one summer to working on the book because I thought this would a good way to get the writing done that I needed. I decided on that means because I thought it would succeed. You decided to enroll in university because you believed it was an effective means to your goal of getting a university education.

Being a critical thinker means that our beliefs should be based on epistemic reasons, and not on emotional or pragmatic ones. Basing one's beliefs on emotions rather than on epistemic reasons is a mistake. Emotions can also make it difficult to collect the evidence we need for our belief to be justified, or even from investigating further. Emotions can also get in the way when we identify too much with our own opinions and beliefs or with our own methods for collecting or evaluating evidence. If I become too emotionally attached to my beliefs and opinions, then I may react negatively when someone asks me for my reasons, or when they raise objections to my belief or when they state their own alternative beliefs. I might feel that they are criticizing me and not just my beliefs. The same is true if I am asked to defend my assessment of the evidence or my use of different methods for collecting evidence. If I come to identify too closely with these particular methods for assessing and collecting evidence, if I come to think of my value as a researcher as tied into their value, then I will react to criticisms of them as if they were criticisms of me and my judgment. This feeling of being under attack might make me feel defensive, and this can prevent me from thinking critically about the issue at hand. The same is true when I ask someone for his or her reasons. This sort of question is easily taken as aggressive or combative, even when the intention is simply to consider the issue from all sides as thoroughly as possible.

BOX 1.2.3C PRACTICAL TIP: DON'T PERSONALIZE REASONS

Reasons and evidence do not belong to anyone; they are **universal**. And whether they are good has nothing to do with who accepts them; they are **objective**. To avoid personalizing reasons, replace the following:

- a. What evidence do you have?
- b. What are your reasons?
- c. Why do you believe that?

with the following impersonal ones:

- a'. What evidence is there?
- b'. What reasons are there to believe that?
- c'. Why should we believe that?

Knowing how to *distance* oneself from one's beliefs and opinions in order to think critically about them is not easy. It is one of the hardest things to achieve. But the best way to avoid this feeling is making sure that one's beliefs and opinions are based on enough of the right kind of evidence. Again: *think twice; decide once*. Another strategy is to avoid talking about "my reasons" or "your reasons" and to talk instead of "the reasons" or "some reasons." This makes sense anyway, since reasons and evidence are not owned or possessed by anyone: they are universal and objective. Instead of asking "What are your reasons for believing that?" which can come across as confrontational, ask, "What reasons are there to believe that?" which makes the question sound less confrontational. Instead of asking, "What is your evidence?" you can ask, "What evidence is there for that?"

Critical thinking requires that we have good epistemic reasons for our beliefs and decisions. Sometimes, in order to decide what to believe or do, we need to *acquire new evidence*. We have several sources of evidence at our disposal, several ways of gaining new information on which to base our decisions about what to believe or what to do. We can gain new evidence through direct observation, testimony, measurement, testing, and experiment. In Chapter 4, we will compare these different sources of evidence and consider when they provide evidence or information that is acceptable. Sometimes, we can decide what to believe or what to do by *drawing conclusions from the evidence we already have*. We can rely on what we already know to compare things or groups of things to see how they are analogous. We can reason about what else has to be true given what we already know or believe that we know. And we can reason about what alternatives the evidence that we have rules out. In chapters 5 and 6, we will compare these different ways of drawing conclusions from the evidence we already have, and study some methods for telling when our reasoning is good.

1.2.4 Good Reasons Are Sufficient and Acceptable

A belief is justified enough for knowledge only if it is based on good enough reasons. Two features are essential to good reasons. First, the reasons have to be *sufficient* to support the belief. Second, the reasons have themselves to be *acceptable*. In later chapters, we will have a lot to say about both these features of good reasons. But let's now take a quick look at each element.

First, a belief is justified enough for knowledge only if it is based on *sufficient* evidence: this just means that it has to be based on *enough* of the right kind of evidence. In deciding what to believe or do we need

to make sure that we have collected enough evidence. This is the idea behind the legal requirement that a jury can find the defendant guilty only if they have proof beyond a reasonable doubt. In deciding whether a defendant is guilty, it is not enough that the prosecution present *some* evidence of guilt. It needs to present enough evidence. Ideally, it should provide enough evidence to guarantee that the verdict the jury reach be the right one. The evidence, in that case, would make it impossible for the verdict to be mistaken. He jury could not go wrong if it made its decision on the basis of that evidence. What is true of juries is just as true of us as we try to decide what to believe and what to do. We can sometimes collect this ideal amount of evidence, but we often have to make do with less than this. In Chapter 3, we will study the ideal amount of evidence and consider some strategies for telling how close we are.

Second, a belief is justified enough for knowledge only if it is based on *acceptable* evidence or information. In a perfect world, we would only rely on evidence that we knew for a fact was true or accurate. But we are rarely in that kind of situation. Usually, we have to make our decisions on the basis of information that we are pretty sure about, but not 100 percent convinced of. Usually, the acceptability of some bit of information of evidence depends on where it came from, on its source. Some sources of evidence are better than others for certain kinds of beliefs, and it is always an important question whether a given source of evidence is trustworthy in a particular case. Direct visual observation is a good source of evidence for beliefs about the colors of objects but it is not a good source of evidence for beliefs about other physical properties of objects. You can often tell just by looking whether something is brown or red, but it is pretty much impossible to tell just by looking whether something will dissolve when placed in water. You can tell by looking whether someone is tall or male, but not whether they are a lawyer or a doctor. You can sometimes tell by looking whether a bridge needs to be repainted but not whether the bridge is at risk of collapse. Some care is needed when we are deciding what to believe or what to do to ensure that our decisions are based on evidence of the right kind.

1.2.4.1 When Evidence Conflicts To make matters even worse, we usually have to make decisions about what to believe or what to do when the evidence we have is in conflict. There are two main ways that our evidence can conflict. Some evidence we have might be **overridden** by other evidence that we have. This happens when the conflicting evidence points in different directions. In a trial, the prosecution might

have circumstantial evidence indicating that the defendant robbed the bank. But if the defense can prove that the defendant was in fact in another country at the time of the robbery, then the prosecution's evidence is overridden. The prosecution's evidence is overridden by the evidence provided by the defense. It would be wrong to rely on the prosecution's evidence in that case, because the other stronger evidence points in the other direction.

To see another example, consider a case of a persistent visual illusion, like the Müller-Lyer Illusion (Fig. 1.1). Every time we look at the drawing, it looks like the black lines are of different lengths. It looks like the middle line is quite a bit longer than either the top or the bottom lines. But if we measure them with a ruler, we will find that they are in fact of the very same length. We are now in a situation where our evidence conflicts. Our eyes tell us one thing; our measurement tells us another. Something has to give. In this case, we have figured out that the evidence we get from direct observation is overridden by evidence we get from the measurement. (Part of what is fun about this illusion is that it is *persistent*: the middle line still looks longer even when we know that it is not.) When we decide what to believe or what to do we have to make sure that we consider all of the evidence we have or can get and we have to make sure that the evidence we decide to go with is not overridden by other evidence. A handy rule of thumb is: Think twice, decide once.

Bits of evidence can conflict in another way. The evidence we have might be **undermined** by other evidence that we have. This happens when we have good reason to think that the evidence we have is from a source that is not trustworthy, either in this particular case or in




Figure 1.1

general. For example, the witness for the prosecution might have testified that he saw the defendant pull the trigger. But if we have good reason to suspect that the witness has a grudge against the defendant, and would lie on the stand, then we might decide not to trust what he says. The evidence that he provided is *undermined* by our reasons for doubting that he is being honest with us. This does not mean that we think that he testified *falsely*. It just means that we should not trust what he says. We should remain agnostic, undecided. We should withhold judgment until more of the evidence is in.

Sometimes we may have reason to question all of the evidence provided by some source. For example, the Müller-Lyer illusion shows that we have to be very careful when we rely on evidence provided by our vision, at least when we are trying to decide when two lines are of the same length. Vision, it seems, can be quite misleading on this kind of topic. But it would be wrong to respond to the Müller-Lyer illusion by believing the opposite of what our eyes tell us. The proper response is to withhold judgment until more evidence is in. “It looks like the middle line is longer, but let’s measure it just to make sure.” When we decide what to believe or what to do we need to make sure that the other evidence that we have does not undermine the evidence we are relying on. Once again: think twice, decide once.

EXERCISE 1.2

1. Short-answer questions:

- a. What is the traditional definition of knowledge?
- b. What is the difference between realism and relativism?
- c. Could one be a realist and a relativist about biology? Why or why not?
- d. Why does the existence of disagreement in some subject area not show that relativism is true of that area?
- e. Why is realism the default attitude to take in a subject area?
- f. How are freedom of action and freedom of belief alike? How are they different?
- g. What is an example of an emotional reason to believe something?
- h. What are emotional reasons not good enough for knowledge?
- i. List two ways in which emotions can be obstacles to critical thinking.

- j. What is the difference between acceptable reasons and sufficient reasons? Give an example of reasons that are sufficient to believe something but not acceptable.
- k. Could evidence be over-ridden without being undermined? Explain using an example.
- l. The traditional philosophical definition of knowledge says that knowledge is justified true belief. When presented with a definition that analyzes some idea or concept into several parts or elements, it is a good idea to investigate how those elements are related to one another. To do this, one asks whether it would be possible to have two of the elements without the third. That is, is it possible for someone to have a belief that is justified (i.e., based on epistemic reasons) even though the belief is not true? Try to construct stories to test whether these elements are independent.

1.3 CRITICAL THINKING AND PERSONAL AUTONOMY

We have seen that critical thinking is aimed at knowledge. It pretty much goes without saying that knowledge is valuable. For one thing, since knowledge requires truth, if we know something then we are not wrong. So critical thinking, to the extent that it can help us gain knowledge, can also save us from making mistakes. And that is a good thing. But thinking critically as we decide what to believe or do is valuable for a different, and in some ways more important, reason. Thinking critically is essential to making up one's own mind, and this is fundamental to being an autonomous person. Let's explore this by looking first at the differences between a belief and a prejudice.

1.3.1 Belief and Prejudice

Knowledge is valuable because of its links to truth. But as we saw, knowledge also requires justification, and justification is valuable because it is what makes the difference between having a well-reasoned belief and having a **prejudice**. And no one wants to be prejudiced. But what exactly is it to be prejudiced, and why is it so bad? Usually when we talk about prejudices we have in mind hurtful views about race, religion, or ethnicity. We say that people who treat Asians or Catholics less well than they treat Europeans or Episcopalians are prejudiced against them. Some people used to believe that Irish immigrants were

lazy, could never keep a job, and didn't care about supporting their families. Irish immigrants were discriminated against as a result of these prejudicial views. Of course, those beliefs about Irish immigrants were totally false, and it is even hard for us now to imagine how anyone could have believed them in the first place. (Unfortunately, it is always easier to spot someone else's prejudices than it is to notice one's own, and there is little reason for optimism that we are any less prejudiced than are the rest of our fellow humans.)

But what made those beliefs prejudices was not that they were false. There has to be a difference between a belief that is false and a belief that is a prejudice. Not every false belief is a prejudice. People used to believe that the Sun orbited the Earth, but we do not think that this false belief made them prejudiced. This was a mistake, not a prejudice. And isn't it possible that a prejudicial view could actually turn out to be true? Suppose I see in the newspaper a picture of someone who has been arrested for some crime and I immediately conclude that he is guilty just from the look on his face. I think we would say that my belief in his guilt was a prejudice. But what if it turned out that he was in fact guilty? Wouldn't we still say that my belief was prejudicial even though it was true? So it seems that whether a belief is a prejudice has nothing to do with whether it is true or false.

One clue to the nature of prejudice comes from the word itself: a prejudice is a prejudgment. To prejudge someone or something is to form a judgment or belief about them before all or enough of the facts are in, before one has enough evidence. Taken literally, a belief is a prejudice when it is not based on good epistemic reasons. My belief in the defendant's guilt was a prejudice because it was not based on good enough reason, and this is so even though the belief was in fact true. The members of a jury are asked not to prejudge the question—not to decide whether the defendant did it—before all of the evidence has been presented. Since critical thinking can help us to make sure that our beliefs are based on good epistemic reasons, critical thinking can also help us to avoid being prejudiced.

1.3.2 Making Up Your Own Mind

But why should we avoid prejudice, especially if prejudice is compatible with being right? This might sound like a silly question, but answering it can help us to see one of the deeper values to critical thinking. One reason to avoid prejudice is that we want to make up our own mind and being prejudiced is the very opposite of making up your own mind. If we let our beliefs get formed before we have had a chance to

examine all the evidence, then it is as if we have lost control of our beliefs and views. Forces and influences that are outside of us would in that case form our beliefs and opinions for us. My instantaneous belief that the arrested man was guilty was not the result of careful deliberation by me. The belief just came over me. I was not responsible for it. In a real sense, I did not *make up* my mind to believe that he was guilty; rather, my mind was made up for me. But that is not how I want my beliefs and opinions to get made. I want them to be *my* beliefs and opinions, ones that I choose and can take responsibility for having, not ones that were simply given to me or (even worse) forced on me. And this is so even if the prejudicial beliefs turn out to be true.

Making up one's own mind is part of what it is to be **autonomous**. Being autonomous means exercising the power to determine one's self, to decide on one's own what to do or what to believe, what kind of life to live. Because critical thinking demands reasons and requires us to be reflective as we decide what to believe and what to do, thinking critically is crucial to exercising our ability to determine our own minds, to decide for ourselves.

Sometimes, making up one's own mind can mean disagreeing with others or even abandoning beliefs and practices that one was born into. Making up one's own mind about religion or politics can sometimes cause pain and lead to separation from those we love. This may be unavoidable if one finds that those practices are not based on good enough reasons, and making the break can require a good deal of courage. But it is not inevitably like that. Examining one's beliefs and practices can also reveal deep and even new reasons for keeping them, and this process can strengthen one's allegiance to them and deepen one's bonds to others who share those beliefs and practices. The benefits of making up one's own mind far outweigh the risks.

We have been discussing the responsibility that we have to make sure that our beliefs and practices are grounded in good reasons. In a famous essay, William Clifford argued that we also have a responsibility to others to make sure that we have good reasons to believe what we do. His reason was that beliefs make a difference to action. We act on our beliefs, and if our beliefs are not based on good enough reasons, then we run the risk that our actions will cause unintended harm. Clifford illustrated this point with a story about a wealthy but penny-pinching ship-owner, whose ship full of immigrants was ready to set sail. The ship-owner had good epistemic reason to think that the ship was not sea-worthy but was reluctant to pay for the needed repairs and to put up with costly delays. Eventually, he convinced himself that the ship was safe. He let his pragmatic and emotional reasons overpower

his epistemic ones. Tragically, he was mistaken and the ship sank, killing everyone on board. Clifford argued that not only was it wrong for the ship-owner to have let the ship sail, it was wrong for him to have believed as he did that it was safe, since his belief was not based on adequate grounds. It is wrong, Clifford insisted, to believe something on the basis of wishful thinking, or for nothing but self-interested reasons.

But suppose that the ship had not sunk. Suppose that the crew and passengers had gotten lucky and the ship made it safely to America. Still, Clifford argued, it would have been just as wrong for the ship-owner to allow the ship to sail and to have believed as he did that the ship was safe. Suppose finally that the ship was in fact quite safe, and that the ship-owner's initial concerns about its safety were not well founded. Still, if the ship-owner ignored those concerns and convinced himself that the ship was safe in hopes of saving a few dollars, Clifford argued, it would still have been just as wrong for him to have allowed the ship to sail and for him to have believed as he did that the ship was safe. It would have been wrong for him to have believed that even though his belief would have been true. It would have been wrong because his belief would have been based on inadequate evidence. It is, Clifford concluded, "Always and everywhere wrong to believe on inadequate evidence." To the extent that critical thinking can help us to ensure that our beliefs are based on strong epistemic reasons it can help us to fulfill the obligations that derive from the ethics of belief.

BOX 1.3.2 CHAPTER SUMMARY

Critical thinking is reasonable and reflective thinking aimed at deciding what to believe or what to do. When we try to decide what to believe or what to do, we are trying to gain knowledge. We want to know the facts or the best way to achieve our goals. Knowledge is justified true belief. A belief is justified only when it is based on enough of the right kind of evidence. Having good epistemic reasons for our beliefs not only reduces the risk of error; it also helps us to avoid prejudice. Avoiding prejudice is essential to making up one's own mind about what to believe and what to do. Critical thinking provides practical methods and standards for helping us to make sure that our beliefs are based on adequate epistemic reasons. In this way, critical thinking helps us to become autonomous.

EXERCISE 1.3**1. Short-answer questions:**

- a. What is the difference between a prejudicial belief and a false belief? Use an example to illustrate your answer.
- b. Could a prejudicial belief be true? Explain, using an example.
- c. Could believing something on the basis of emotions alone make one autonomous? Why or why not? Use an example to illustrate your answer.
- d. List five character traits that you think are characteristic of an ideal critical thinker. Give an example of each one.
- e. Suppose that Jones is a universal relativist (i.e., a relativist about all subject matters) and that Smith is a universal realist (i.e., a realist about all subject matters). Could they nonetheless agree on all the facts? What exactly would they disagree about?
- f. Consider the following proposed definition: to lie is to deliberately say something that is false and that one believes to be false in order to mislead another person. What are the elements of this proposed definition? Use the Test for Conceptual Independence to determine whether the elements are independent of one another.
- g. Some philosophers claim that it is wrong to lie to someone because it prevents them from making up their own mind. Construct a story about Jones (or your favorite character) to illustrate this point. Do you agree that this is part of what makes lying wrong?
- h. Thomas Jefferson is supposed to have said that all knowledge begins with book knowledge; that is, with knowledge that we get from reading books or from trusting what other people say. Could this be right?

2. In the following texts, reasons are given for some belief or practice. Explain whether the reasons are epistemic ones.

- a. Sally believes that it is wrong to eat meat. She once watched a documentary on the methods used to kill cows, and it made her so sad that she immediately became opposed to eating meat.
- b. The glass of milk is empty. I can see with my own eyes that it is.
- c. The glass of milk is empty. I can't see it, but my mother just told me that it is.

- d. We have to hold the party on Christmas Eve, because we have always held it then.
 - e. John believes that the sofa will fit up the stairs. He first measured the sofa and then the stairs, and decided that it would go up easily if tilted on its side.
 - f. Ashanti believes that Senator Doolittle's proposal is not cost effective. She finds that politicians are such hypocrites that she disagrees with everything they propose.
 - g. Robert believes that his car will not last much longer. He knows several people who own the same make of car and none of them lasted as long as his has lasted. So he figures that his car will not last much longer.
 - h. Susan believes that birds are a kind of dinosaur. She does not remember how she first came to believe it, but has decided to believe it until she finds some contrary evidence.
 - i. John thinks that smoking causes cancer. He believes it because his mother and two aunts died of cancer after smoking all of their lives.
3. In each of the following, several epistemic reasons are given to believe something. Which is the strongest reason? What makes it stronger?
- a. John, Susan, and Terry all believe that the bank robber was a male. John was there during the robbery and saw the robber. Susan read about the robbery in the newspaper. Susan told Terry about the robbery.
 - b. John and Susan both believe that the acid caused the chemical reaction. John read in a textbook about the likely causes of such a reaction. Susan performed several experiments to rule out other possible causes.
 - c. Susan and Terry both believe that their checking accounts are overdrawn. Terry got a phone call from his bank telling him about his balance. Susan noticed it when she was balancing her checkbook last night.
 - d. John and Susan believe that some early settlers in New England suffered real hardships. John read some original diaries written by early settlers. Susan saw a documentary on TV.
 - e. John and Susan both believe that building a new bridge will greatly reduce the current traffic problems. John based his belief on a comparison of the proposed bridge and the traffic problems to those in other cities. Susan believes it because she heard the

- city planners claim that the bridge would reduce traffic problems.
- f. John and Susan both believe that raising the minimum wage would lead to higher unemployment among the very poor. John believes it because he thinks that it follows from what he learned in his economics class. Susan believes it because she works in an unemployment office and has seen the unemployment lines grow after the wage has been raised in the past.
4. In (a) in exercise (3), if the belief had been that the robber was a male with a long criminal record, then Susan's belief would have been better justified than John's, since it is hard to tell just by looking whether someone has a criminal record, but this is the kind of information a newspaper report would get right. For each of the other questions in (3), change the shared belief but not the kind of evidence each character relied on so that the other person's reasons are stronger.

1.4 MISTAKES TO AVOID

This book is intended as a practical guide to deciding what to believe. In later chapters we will discuss some strategies and standards that can help us to make sure that our decisions about what to believe or do are based on good epistemic reasons. As we go along, we'll draw attention to some familiar mistakes, sometimes called "fallacies." Identifying them will help us to avoid them in our own thinking and to spot them in other people's thinking. Seeing why they are mistakes will help us know what to look for as we try to find good epistemic reasons for our decisions. All of the mistakes are collected together at the end of the book, for quick and easy reference.

Appeal to Origins It is a mistake to assume that a belief's originating reasons are epistemic reasons too. There are many factors that influence what one believes, and not all of them need be epistemic. Jones' belief that humans evolved from other species might have been originally based on excellent reasons he learned in school. In that case, the originating reasons would be epistemic ones. But maybe he took on that belief as an act of rebellion against his parents. In that case, his original reasons for believing that are not at all epistemic ones. It is a mistake to assume that a belief's originating reasons are epistemic reasons too.

Personalizing Reasons It is a mistake to personalize reasons by treating them as if they belonged to someone. Epistemic reasons are universal: if they are reasons for me to believe something then they are equally reasons for anyone else to believe that thing. Reasons are also objective: whether they are good reasons has nothing to do with me or with anyone else. Personalizing reasons can obscure the fact that they are universal and objective. It can also allow emotion to get in the way of thinking critically if one identifies too much with one's own reasons or if one rejects reasons just because someone else accepts them.

Appeal to Relativism It is a mistake to assume that truth is relative. Relativism with respect to some subject matter is the view that the facts in that area are in some way dependent on our beliefs about them. We noted in Section 1.2.1 that relativism might be the right attitude to take towards such topics as what is humorous or what is tasty. But for most topics, even religious and moral ones, it is best to assume that realism is the appropriate attitude, unless one has powerful reasons not to. For most topics, in other words, it is wrong to assume that what is true for me might not be true for you, or that what is true for our community or culture might not be true for others. Truth is the same for everyone.

Sometimes, an appeal to relativism will be used as an attempt to bring a discussion to an end. One person, perhaps tired of the debate or feeling that they are on the losing side, will say to the others: "Well, I'm entitled to my view and you are entitled to yours." This kind of response is fine if what is intended is that everyone is allowed to make up their own minds about what to believe or do. But if the point is that we can both be right even when we disagree, then this is a mistake that we should avoid, unless there is excellent reason to think otherwise.

Appeal to Emotion It is a mistake to base our beliefs only on our emotions. For a belief to be justified enough for knowledge it must be based on good epistemic reasons. Epistemic reasons are reasons to think that the belief is true. Emotional reasons are not epistemic ones. How a belief makes us feel has nothing to do with whether the belief is true. As we have already noted, critical thinking does not aim to eliminate emotion from our decision-making. I doubt this would be possible even if it were worthwhile. Many of our beliefs are so fundamental to our deepest conceptions of ourselves, of our culture or our place in the universe that the pain involved in abandoning them would be too great to bear. It is fine for our beliefs to have or even constitute these emotional supports, so long as they also have sufficient support

from epistemic reasons. But it is a mistake to base our beliefs on nothing but emotional reasons. We also saw that it is a mistake to allow emotions to prevent us from collecting or assessing the evidence we need to make the decisions we must.

Privileging Available Evidence It is a mistake to assume that evidence that we currently have is more acceptable or more sufficient than evidence that we might collect. It might be that our current evidence is the best we can get. But we will not know this until we try to collect more. Crucially, even if we have excellent reason to rest content with the evidence we have, we should always keep an open mind that we might uncover new evidence that will over-ride or undermine the evidence we now have.

BOX 1.4 MISTAKES TO AVOID: APPEALING TO TRADITION

It is a mistake to rely on some method for solving a problem, or to adopt some standard of evidence, just because it is the traditional method or standard. Critical thinking is reflective in that it requires thinking about and evaluating these methods and standards themselves. Thinking outside the box can involve experimenting with new methods and standards, as well as considering new possibilities. A story from my own life nicely illustrates this point. When I was a child, my family lived in England. The houses in our neighborhood all had their water pipes running up the outside of the house, instead of inside the exterior walls. Predictably, the pipes froze and burst every winter and workmen had to be called to repair them. My father asked the landlord why the pipes were on the outside instead of inside the walls. The landlord explained: Well, if they were on the inside, then we could not get to them when they froze. The landlord had accepted the traditional way of thinking of the problem: he saw it as an access problem, best solved by putting the pipes on the outside walls. Having been raised in Canada, my father saw it as a freezing problem, best solved by putting the pipes inside the heated space of the house. The landlord's mistake was in not asking whether the traditional way of thinking was the right way.

Appeal to Tradition It is a mistake to believe or do something simply because that belief or practice is traditional. The fact that a belief or practice has a long history is not an epistemic reason to continue it. Its

history cannot show that the belief is true or that the practice is worthwhile. But as with appeals to emotion, the point is not that we should work to avoid relying on tradition. Indeed, we probably could not know everything we do if we did not rely on others. In Chapter 4, we will consider when we can trust the evidence provided by other people including our ancestors. We will see that it is not that hard to decide when we have good epistemic reasons to believe what they tell us. But it is a mistake to rely on what others tell us without also relying on epistemic reasons.

1.5 PRACTICAL STRATEGIES

This book is intended as a *practical* guide to deciding what to believe. As we go along, we'll draw attention to some useful practical strategies or methods. These will all be collected into an appendix at the end, for quick and easy reference.

Testing for Conceptual Independence It is good to know how to test for conceptual independence. Twice so far we have considered definitions. The first was Robert Ennis' definition of critical thinking as reasonable, reflective thinking aimed at deciding what to believe or what to. The other was the standard philosophical definition of knowledge as justified, true belief. When an idea or concept is analyzed into several parts or elements, it is always a good idea to ask whether those parts or elements are conceptually independent of one another. To do this, simply ask yourself whether you can think of an example of something that has some of the elements but not others. For instance, we noticed that simple arithmetical calculations are a kind of thinking aimed at deciding what to believe but are not reflective because they do not require thinking about the method one uses. This example shows that reflective thinking and thinking that is aimed at deciding what to believe or do are conceptually independent. In one of the chapter's exercises, you discovered that in the case of knowledge, a belief's being true is conceptually independent of its being justified. Whenever a concept or problem has elements or parts, ask: how are those parts related to one another?

Think Twice; Decide Once To paraphrase the old carpenter's motto (measure twice and cut once), it is best to think twice and decide once. We know from psychological experiments that people are reluctant to change their minds. Once our opinions are set, it seems to take a lot

of doing to revise them. For one thing, people tend to privilege evidence that confirms their already existing beliefs over evidence that conflicts with it. They assume that evidence that conflicts with what they already believe is probably not reliable. For another thing, people tend to prefer the evidence they have to evidence they would have to do something to get. To protect against these built-in obstacles to critical thinking, it is better to make sure that one has enough of the right kind of evidence before one makes a decision. It is better to think twice and decide once than to have to go back and revise one's decisions.

1.6 FROM THEORY TO PRACTICE: APPLYING WHAT WE HAVE LEARNED

One goal of this book is to provide you with the conceptual tools and the practical strategies you need to become a strong critical thinker. Thinking critically requires having an appropriate vocabulary for describing and evaluating the decisions we need to make, as well as having the strategies and methods needed to make sure that our decisions are based on the right kind of evidence. But book learning only goes so far. Becoming a critical thinker requires using these concepts and skills in our own life. We can and should think critically about our own decisions and values, about our classes and studies and about our workplace experiences. The following set of exercises will continue throughout the book, as we acquire new concepts and learn new strategies. They are designed to help you “transfer” what you learn in this class to the rest of your life. And as with everything, the more you put into them, the more you will get out.

Thinking Critically about Ourselves Good critical thinking begins at home. This means that we can practice the skills and strategies involved in thinking critically by reflecting on ourselves and our own decisions and values. The self-examination exercise—which continues throughout the book—asks you to examine your conception of a good person. In this chapter, we will begin by outlining the exercise:

- a. List five or six traits that you think are essential to being a morally good person. You can be as specific or as general as you like. But it is good to pick traits that are as varied as you can. Some examples: honesty, loyalty, generosity, and faithfulness.
- b. Pick one of them to work on for the remainder of the text. Try to define it in other terms, as if you were explaining it to someone

who was unfamiliar with it. Think up a story in which it is illustrated.

- c. Explain why you think that trait is essential to being a morally good person. Try to make sure that your reasons are epistemic ones, as opposed to emotional or pragmatic ones.

Thinking Critically in the Classroom Every university and college is in the business of producing critical thinkers, and each of their departments and programs are charged with trying to improve the critical thinking skills of its students. Geology departments want to do more than just teach their students geological facts; they want to teach them how to think critically about geology. Business programs want to help their students become adept at thinking critically about business problems and solutions, and not just to teach them business concepts and practices. This exercise, which will continue throughout the text, is designed to help you see where critical thinking can fit in with your studies:

- a. In your own words, and with as much detail as you can, list five or six things in your program where critical thinking is required in learning. Some examples: memorizing definitions and concepts; learning historical events and explanations; performing measurements; collecting evidence; doing factual research; writing essays; performing experiments; evaluating performances and works of art; analyzing texts and arguments. As clearly as you can, and using the concepts we have studied in this chapter, explain in what way critical thinking is required in each of them.
- b. Using the textbooks for your courses as a guide, compile a list of the five or six most important concepts for your field of study. These will be the concepts that are used most broadly to formulate the claims and to frame the subject matter. They should not be the same as the concepts in another field. For example, the concept of a cell is essential to biology, but not to economics; the concept of demand is crucial for economics, but not to history.
- c. The only way to succeed in your studies is to study hard. Critical thinking can help with this. List five or six things that you do as part of your studying and describe how they involve critical thinking, in the sense that we have been discussing in this chapter. Some might involve decision making while others involve reflection.

Thinking Critically at Work Studies show that employers value an ability to think critically more than just about any other trait in an employee. They want their workers to be able to think critically about both day-to-day problems as well as about broader organizational performance and plans. Many employers even provide critical thinking training as an element in management development. This exercise, which will continue throughout the text, is designed to help you see where critical thinking can be applied at work:

- a. Thinking about your workplace, list five or six tasks that you or your co-workers are regularly asked to perform that require thinking critically in the sense that we have been discussing. They can be as simple or complex as you like, but again it is best to make the list as varied and specific as possible. (A hint: start with very general tasks, and then analyze them down into smaller more discrete tasks.) Some examples: dealing with customer complaints; regular communicating with co-workers and supervisors; ordering and stocking inventory; dealing with late or delinquent bills; implementing or evaluating systems and procedures.
- b. Pick one of those tasks, and answer the following.
 - i. What is the task? Be as detailed and specific as you can.
 - ii. In what ways does it require critical thinking? Which of the elements of critical thinking does it require?
 - iii. What information do you usually need to perform the task and how do you usually collect and assess that information?
 - iv. If you could implement a change that would improve or enhance your performance of that task, what would it be?
 - v. What obstacles are there to thinking critically in the performance of that task? Be as specific and detailed as you can.

This page intentionally left blank

2

CLARIFYING MEANING

Critical thinking is reasonable, reflective thinking aimed at deciding what to believe and what to do. Knowing how to reflect critically on meaning is fundamental to critical thinking. Before we decide what to believe or what to do, we need to make sure that we have clearly defined the words and concepts that we use to formulate the statements that we are assessing, to describe the proposals we are considering, and to frame the problems we are facing or the solutions we are contemplating. Otherwise, we run the real risk that we'll end up believing something we shouldn't, doing something that won't succeed, or failing to solve the problems we tackle. In this chapter, we will study some practical strategies for constructing and evaluating definitions. But let's start by discussing in a bit more detail where definitions fit into critical thinking.

2.1 THE PLACE OF DEFINITIONS IN CRITICAL THINKING

First, knowing how to construct and evaluate definitions is fundamental to attaining the goals of critical thinking itself. If we are trying to decide whether to accept or believe some claim or statement, then we need to

make sure that we fully understand what the claim or statement means. We may need to analyze it into elements, contrast and compare it with similar claims, and determine what else accepting or believing it would commit us to. If we are presented with several proposals or plans of action, several options for reaching some end, then before we decide on one of them we need to make sure that we understand how the plans differ and how they are similar. What would be involved in adopting one of them over the others? What other possible courses of action are there and what would be involved in adopting them instead? Knowing how to clarify meaning by reflecting critically on it is thus fundamental to being a critical thinker.

Second, knowing how to construct and evaluate definitions is fundamental to understanding and fully engaging in an academic discipline. Every discipline or field of study has its own **fundamental concepts**, ideas, and technical terms. The concepts that are basic to biology, for instance, are very different from those that are basic to chemistry or physics. The concepts and ideas that are characteristic of archaeology are very different from those that are fundamental to anthropology or to the literary analysis of drama. These fields of study all have different fundamental concepts even though they may study the very same phenomena. Biology, chemistry, and physics all study living systems (among other things), just as archeologists, anthropologists, and literary theorists might all study ancient Greek tragedy. Different disciplines approach the same phenomena from different perspectives, and these perspectives are defined—or **framed**—by the concepts that are fundamental to each discipline. This difference is part of what makes thinking about living things from a “biological” perspective so different from thinking about them from a physical or sociological one. The same is true of business organizations. Each organization has its own ways of describing its structure and its operations and goals. Participating successfully in an organization requires understanding these concepts, and this requires knowing how to think critically about meaning. Knowing how to critically reflect on the meaning of a discipline’s or an organization’s fundamental concepts is thus fundamental to being able to participate in that discipline or organization.

Finally, knowing how to construct and evaluate a definition is often a fundamental step in solving problems. Sometimes, knowing how to frame or define a problem is half the work needed to solve it. Some problems are so clearly defined from the start that little or no critical reflection is required to solve them. Simple arithmetical calculations are like this. We know up front what the problem is (e.g., find the

square root of some number), we know what methods are to be used, and we know what kind of answer we are looking for. But for other kinds of problems, it is not clear at the start just what the problem is, or what the best method is for approaching or solving it. With such **open-ended** problems, it might not even be clear what will count as an acceptable solution, and to solve them we may need to think hard about how to formulate the problem. The problems of ending poverty, of rejuvenating American cities, or of designing a university's curriculum to improve student performance are open-ended problems, not because they cannot be solved, but because part of the problem is getting clear on what the problem is. Knowing how to reflect critically on problems with the goal of defining them clearly is thus central to problem solving.

In this chapter, we will study what is involved in reflecting critically on meaning. We will study a practical strategy for constructing and evaluating definitions, whether of concepts, plans, or problems. We will see how this method can be used to understand the fundamental concepts that disciplines use to frame the phenomena they study.

2.2 ASSERTION

It will be helpful to start by drawing attention to some basic facts about language, since we use language to formulate our claims and proposals. Some of the facts we will discuss may seem quite straightforward, but that is okay, since others of them are among the more difficult concepts in this whole book. Altogether, discussing them will provide a secure foundation for our later discussions about definitions and fundamental concepts.

Inevitably, we use language when we try to decide what to believe and what to do. We use language to express our beliefs, to state our evidence, and to present our reasons. Our native language is so familiar to us that it can be hard to appreciate just how complex it really is, or just how much the clarity of our reasoning depends on the care with which we use our language. In this section, we will study some of the complexities involved in using language to reason.

Let us start with the basics. To **assert** something is simply to claim that it is true. If I assert that oysters are delicious, or that it is raining out, or that Barack Obama is president, then I am claiming that it is true that oysters are delicious, or that it is raining out, or that Obama is president. Anything that can be true or false is something that can be asserted, though we might lack the words for it. Our

language has tremendous expressive power; we can say a lot with it. But we cannot say everything there is to say. We assert something when we say what we believe. Of course, one can also assert something without believing it; this is what makes lying possible. And we assert things when we provide our reasons for our beliefs too. When we state our reasons for believing something, when we collect together our evidence in support of a point of view or proposal, we are making assertions.

2.2.1 Propositions

Though we use sentences to make assertions, we do not assert sentences. We do not assert the sentence “Oysters are delicious”; rather, we use assert the **proposition** that oysters are delicious. A proposition is simply a bit of possible information, something that can be true or false. The proposition that oysters are delicious is something that we can claim to be true or believe to be true. It is something that we can doubt or question. It is something we can clarify and analyze. Most importantly for our present purposes, a proposition is something that we can assert and argue for or against.

There are several linguistic points to keep in mind about assertion and propositions. First, to assert something I need a sentence, but not any sentence will do. Here is a partial list of sentences that cannot be used to assert something:

Are oysters delicious?
 I baptize you, “David.”
 I promise to pay you back next week.
 Stop!

These are perfectly fine sentences and have their uses, but they cannot be used to claim that something is or is not the case. They are used to ask something, to perform a baptism, to make a promise, or to issue a command. But they cannot be used to state truths. There are interesting and important things to say about how to use these kinds of sentences. But our interest in this book is with reasoning that is aimed at truth. So we will focus on assertions, instead of commands, questions, and other speech acts.

Second, different sentences can be used to assert the very same proposition. If I want to assert that John is a liar, I could do it using either of the following sentences:

John is a liar.

John is mendacious.

These sentences are simply different ways of putting the same proposition into words. They are synonymous. It is a good thing that different sentences can be used to assert the very same proposition, since otherwise people who spoke different languages could never say the same things! But it does make our work as critical thinkers a little bit more difficult. For it means that we cannot keep track of what someone is asserting just by keeping track of what sentences they are using. We need to keep an eye on the possibility that they are simply repeating themselves in other words.

The third point is a related one: a single sentence can be used to assert more than one proposition. Suppose Jones used the following sentence to make an assertion:

Bill's wife is mad at him.

If we wanted to figure out just what propositions Jones had asserted, we might come up with the following list:

Bill has a wife.

Bill's wife is mad.

Bill's wife is mad at him.

Each of these things is on our list because each of them is a proposition that Jones was claiming to be true in asserting what he did. In asserting what he did, Jones in fact asserted several propositions. Of course, rather than use three sentences, Jones did what we would all do and used a single sentence to assert all of them at once. This is another handy trick made possible by language. But it too makes our work as critical thinkers a bit more difficult, since (as before) it means that keeping track of what someone is asserting is not the same as keeping track of what sentences he is using. If we want to make sure that we know what someone is asserting and—what is just as important—what they are not asserting, we need to know how to analyze their assertions into their constituent propositions. After all, if they are giving us reasons to believe something, then each proposition might be a separate reason, even if they are all packed neatly into one long sentence. If we do not distinguish one from the other, then we will not

really appreciate his reasons. If we are not careful, we may disagree with something that was not even asserted! We will return to this important point in a later chapter.

EXERCISE 2.2.1

1. Short-answer questions:

- a. What is it to assert something?
- b. Is a prediction of some future event an assertion? Explain using an example.
- c. Could the word *Yes* be used all on its own to make an assertion? If so, use an example to illustrate.
- d. Could a false sentence be used to make an assertion? Use an example to illustrate.
- e. Do you have to believe what you assert? Explain using an example.
- f. What is the difference between a sentence and a proposition?
- g. Compose a sentence that asserts more than three propositions.
- h. Compose a sentence that asserts only one proposition.
- i. How many propositions are asserted in the following:
 - i. The Earth orbits the Sun quickly.
 - ii. Can I have another cookie?
 - iii. I think that cookies are delicious.
 - iv. I love you!

2.2.2 Assertion Test

Here is a test we can use to tell whether a given proposition is one of the propositions being asserted. Let's call it the **Assertion Test**. Simply ask: *could the assertion be true even if that proposition were false*. If the answer is yes, then that proposition was not asserted. The idea behind the test is this. To assert a proposition is to claim that it is true. So if an assertion could be true even if a given proposition were false, then that proposition was not asserted. It could not be part of what is being asserted whether the assertion is true does not depend on whether it is true. On the other hand, if the answer is no—that is, if that assertion would be false if that proposition were false—then that proposition was one of the things that was asserted.

BOX 2.2.2A PRACTICAL TIP: THE ASSERTION TEST

To tell whether some proposition is asserted in a sentence, ask whether the sentence as a whole could be true even if that proposition were false:

If Yes, then that proposition is not asserted.

If No, then that proposition is asserted.

The Assertion Test will prove enormously useful not only with definitions but in later chapters too. But it is a bit tricky to use, at least at first. Let's try it out on a relatively easy case. Consider the following sentence:

Stephen Harper is prime minister of Canada and Barack Obama is the President of the United States.

This sentence is a **conjunction**. That means that it is a sentence containing the word "and." We can identify two propositions (or *conjuncts*) joined by that "and":

- i. Stephen Harper is prime minister of Canada
- ii. Barack Obama is president of the United States.

Let's use the Assertion Test to see whether these conjuncts are asserted in that sentence. First, suppose that (i) were false; that is, suppose that Stephen Harper was not the prime minister of Canada. That is a bit hard to do, since he is in fact prime minister. But we can all imagine that the election might have gone differently in such a way that he had not been elected. Now that we are supposing that it is false, let's ask whether that long sentence of which (i) is a part could still be true. That is, if Stephen Harper were not the prime minister of Canada, could it still be true that Stephen Harper was the prime minister of Canada and Barack Obama was president of the United States? Obviously, it could not be true. If Stephen Harper were not prime minister of Canada, then there is no way that that longer sentence could be used to say something true. For one of the things one would be asserting in using that sentence is precisely that Stephen Harper is the prime minister of Canada. The same holds for the other conjunct. It too is

asserted in using that sentence. The bottom line, then, is this: in asserting a conjunction, both conjuncts are asserted. (Another way of putting that point is this: a conjunction is true if, but only if, both conjuncts are true.)

BOX 2.2.2B SUMMARY: ASSERTING A CONJUNCTION

In a conjunction, both conjuncts are asserted. “A and B” asserts that A is true and that B is true. It is thus false if either conjunct is false.

BOX 2.2.2C SUMMARY: ASSERTING A DISJUNCTION

In a disjunction, neither disjunct is asserted.
 “A or B” does not assert that A and it does not assert that B.

That case was pretty straightforward. Here’s a more difficult one:

Either Stephen Harper is prime minister of Canada or Barack Obama is president of the United States.

This sentence is a **disjunction**, which just means that it is a sentence containing the word “or.” We can again identify the same two propositions, though because in this case they are joined by an “or,” we call them *disjuncts*. Let’s use the Assertion Test to see whether the disjuncts are asserted. Is the proposition that Harper is prime minister of Canada asserted in that disjunction? Could the entire disjunction be true even if that disjunct were false? The answer is yes. All it takes for that disjunction to be true is for *one or the other* of its disjuncts to be true. *They do not both need to be true.* The same goes, of course, for the other disjunct. So, in asserting a disjunction, neither disjunct is asserted. Instead, in asserting a disjunction one is asserting that one or the other (or perhaps both) of the disjuncts is true, but that is not the same as asserting either disjunct.

Now consider the following **conditional**:

If Stephen Harper is prime minister of Canada, then the prime minister of Canada is male.

A conditional is just a sentence containing an “if, then.” The proposition following the “if” is the *antecedent* (in this case, the antecedent is

the proposition that Stephen Harper is prime minister of Canada.) The proposition following the “then” is the *consequent* (in this case, the consequent is the proposition that the prime minister of Canada is male.) Let’s use the Assertion Test to see whether either the antecedent or the conditional is asserted. Suppose that Stephen Harper was not the prime minister of Canada. Could it still be true that if he had been, then the prime minister of Canada would have been male? Yes, that would still be true. Now consider the consequent. Suppose that the prime minister of Canada was not male. Would it still be true that if Harper had been prime minister then the prime minister of Canada would in that case have been male? Again, the answer is yes. This shows that in a conditional assertion, neither the antecedent nor the consequent is asserted. Rather, in asserting a conditional one is asserting that a complex logical relation holds between the antecedent and the consequent; roughly, that if the antecedent were true, the consequent would be true too. Conditionals play an essential role in our reasoning using definitions and in our reasoning about causal relations among events (which we will study in Chapter 5).

BOX 2.2.2D SUMMARY: ASSERTING A CONDITIONAL

In a conditional, neither the antecedent nor the consequent is asserted. “If A, then B” does not assert that A is true or that B is true. Rather, it asserts that the truth of A is sufficient for the truth of B.

So far, we have seen that disjunctions and conditionals do not assert all of the propositions we can identify in them. There is a final case that is worth noting. Consider the following two sentences:

Joan believes (says) that Stephen Harper is prime minister of Canada.
Joan knows that Stephen Harper is prime minister of Canada.

Both contain a **noun clause**, which is just a sentence starting with the word “that.” The proposition in that noun clause is the same in both cases, namely that Stephen Harper is the prime minister of Canada. Is it asserted in either sentence? Again, we can use our Assertion Test to answer the question. Suppose that Stephen Harper was not the prime minister of Canada. Could it still be true that Joan believes (or says) that he is? Yes. She might believe it or say it even if it was not true. We all know from personal experience that it is possible to speak or

believe false things, though we all try to avoid this. So in using this sentence, one would not be asserting that Stephen Harper is prime minister of Canada. What about the second sentence? Could it be true that Jane *knows* that Harper is prime minister of Canada even if, as we are supposing, he was not prime minister of Canada? Here the answer has got to be no. If you know something, then what you know has got to be true. So in using that second sentence, one would be asserting that Stephen Harper is prime minister of Canada. In the first sentence, that proposition is part of a **nonasserted noun clause**, whereas in the second sentence it is part of an **asserted noun clause**.

BOX 2.2.2E SUMMARY: ASSERTED AND NONASSERTED NOUN CLAUSES

Some sentences containing a noun clause assert the proposition in the noun-clause and others do not.

These are nonasserted noun clauses:

“believes that,” “hopes that,” “thinks that,” “says that,” “asserts that”

These are examples of asserted noun-clauses:

“knows that,” “remembers that,” “proves that,” “sees that”

The Assertion Test is a bit tricky to use. It requires supposing something to be false, and then asking whether something else could still be true. This requires imagination and flexibility. But knowing how to tell what is being asserted and what is not is fundamental to critical thinking.

EXERCISE 2.2.2

1. Short-answer questions:
 - a. What is an assertion?
 - b. What is a conditional?
 - c. What is the difference between a disjunction and a conjunction?
 - d. What is a nonasserted noun clause?

- e. Can you use a disjunction to make an assertion? Give an example.
 - f. Could a disjunction be true if both of its disjuncts were false?
 - g. Could a conjunction be true if one of its disjuncts was false?
 - h. Using the concepts we have been discussing, explain why there is something wrong with using the following sentence to make an assertion: “It is now raining, but I do not believe that it is now raining.”
2. What is asserted in the following?
 - a. Susan sold Peter the chips and he ate them all.
 - b. Either the cook is angry or her spices are hotter than usual.
 - c. The restaurant was pretty disappointing: the fish was overcooked, the sauces were boring, and the service was terrible.
 - d. Padma told me that her mother is coming to town.
 - e. Harry knows that Voldemort is back.
 - f. The First City Bank has gone bankrupt and will be closing its doors next Monday, according to bank officials.
 - g. The United Nations military intervention in Africa has produced more harm than good and has not achieved any of its main objectives, reports a local think tank.
 3. Construct three sentences in which the following is expressed but not asserted:
 - a. Oysters are delicious.
 - b. Sam believes that oysters are delicious.
 - c. Johannes knows that oysters are delicious.
 4. Look at today’s newspaper and find three letters to the editor. List all of the assertions made in each one.

2.3 CONSTRUCTING AND EVALUATING DEFINITIONS

When we think of definitions we probably think first of the definitions of words or concepts. But we can also define plans and problems. To define something is just to make it clear, to distinguish it from other things with which it might be easily confused, and we can make just about anything clear. As we noted at the outset of this chapter, knowing how to construct and evaluate a definition is fundamental to critical thinking. Whether we are trying to decide whether to believe some assertion, whether to adopt some course of action, or how best to solve some problem, we have to make sure that we have a clear

understanding of what the claim, course of action, or problem is. And this means that we have to know how to define it. In this section we will study a practical, four-step method for constructing a definition and we will apply it in several cases.

BOX 2.3 CRITICAL THINKING IN CRIMINAL JUSTICE

How many child kidnappings would you say there are every year in all of the United States? According to the National Criminal Information Center (NCIC) database, over 800,000! That is 2,000 kidnappings every week. That is more than the population of many American towns. Indeed, it is more than the entire population of Vermont, Wyoming, South Dakota, North Dakota, and Alaska. But what do they mean by “child kidnapping”? It turns out that, according to the NCIC definition, child kidnapping includes: abductions by a family member, abductions by someone other than a family member, runaways, abandoned children, and lost or otherwise missing children. Most of these would not ordinarily be counted as kidnappings. In fact, of the cases that the NCIC included as child kidnappings, only 115 were stereotypical kidnappings: a nonfamily abduction perpetrated by a slight acquaintance or stranger in which a child is detained overnight, transported at least 50 miles, held for ransom, or abducted with the intent to keep the child permanently or killed. This puts a very different face on the facts. The point here is not that the 800,000 cases are not serious or troubling or even tragic, but only that if we do not know what is meant by “child kidnapping” we will not know what to make of that number.

Here is a four-step method (the “SEEC” method) for constructing and evaluating definitions.² A definition should state the meaning as clearly as possible and in as few words as possible. This statement can usually take the form of a **slogan**. A definition should **expand** on that statement by filling in some of the detail that inevitably will get left out of a succinct statement. Among other things, the elaboration might say something about how the different elements in the statement are related one to another. This should take no more than a few sentences. A definition should provide an **example** or two, depending on the

²This method is inspired, in part, by Richard Paul’s *Critical Thinking: Basic Theory and Instructional Structures* (Rohnert Park, CA: Foundation for Critical Thinking, 1998).

complexity of what is being defined. The example could be from real life or it could be fictional, so long as it is clear and uncontroversial. Finally, a definition should identify some **contrasting** cases with which the thing being defined might easily be confused.

2.3.1 Slogan

Ideally, the statement of a word's meaning should say what it takes for something to be a **referent** of the word. The referents of a word are just whatever the word is a word for. The referents of "apple pie" are just all the apple pies there are, or ever have been, or ever could be. Likewise, the referents of "kidnapping" are just all the actual and possible kidnappings. A definition of "apple pie" should state what it is for something to be the referent of "apple pie." In effect, this comes to the same as saying what it is for something to be an apple pie. We can be a bit more precise about what this demand comes to.

2.3.1.1 Necessary and Sufficient Conditions It would be good for a definition of a word to state **necessary** and **sufficient** conditions for being a referent of the word that everyone would agree on. A condition is necessary when something has to have it to be one of the word's referents. For instance, containing apples is necessary for being an apple pie. Nothing has been, will be, or ever could count as an apple pie unless it had some apples in it, whether sliced, diced, or pureed. Containing apples just is part of what it is for something to be an apple pie. Arguably, having a bottom crust is also necessary. Perhaps having a top crust is also necessary.

A condition or set of conditions is sufficient if that is all it takes to be a referent of the word. Is containing apples sufficient for being an apple pie? No, since an apple-filled donut contains apples but is not an apple pie. Is having a top and bottom crust and containing apples enough to make something an apple pie? Could something have those three feature or conditions and still not qualify as an apple pie? Some might prefer apple pies that contain sugar or cinnamon, but it is not obvious that those are really needed. They might be needed for something to be a tasty or good apple pie. But I don't think they are needed for something just to be an apple pie.

2.3.1.2 Necessary and Sufficient Conditions Are Related in Complex Ways A condition that is necessary for being a referent of some word does not have to be sufficient. Containing apples is necessary but not sufficient for being an apple pie. Conversely, conditions that are

sufficient for being a referent of some word do not have to be necessary as well. Consider the concept of being a millionaire. Having at least one million dollars is necessary, since every possible millionaire must have at least that much. Having 100 million dollars is sufficient for being a millionaire: anyone with that much money is a millionaire. But having that much money is not necessary, even though it is sufficient. Likewise, getting an A is sufficient for passing a course, but it is not necessary. So conditions that are sufficient for being the referent of a word need not be necessary as well.

We can use the **Assertion Test** to tell whether a proposed condition really is necessary or sufficient. Suppose that we want to define the concept *murder*, and want to know whether being a *deliberate killing* is necessary or sufficient for being a murder. If being a deliberate killing is *necessary* for being a murder, then asserting that something is a murder is also asserting that it is a deliberate killing. So to tell whether it is necessary, we can ask the following: if we assert that something is a murder are we also asserting that it is a deliberate killing? I think the answer to this is yes. Any case of murder would also be a case of deliberate killing. This shows that being a deliberate killing is necessary to being a murder. Is it also sufficient? If it is, then asserting that something is a deliberate killing would also be to assert that it is a murder. So, to tell whether it is sufficient we ask: if we assert that something is a deliberate killing, are we also asserting that it is a murder? Here, I think, the answer is no. Killing someone in self-defense might be a deliberate killing, but would not be murder. So being a deliberate killing is not sufficient for being a murder.

2.3.1.3 Narrow and Broad Definitions A proposed definition is too **narrow** if it includes conditions that are not really necessary. The following proposed definition is too narrow: An apple pie is a pie-shaped, apple-filled, and cinnamon-flavored pastry. This is too narrow because it excludes things that really are apple pies, such as the rectangular apple pies sometimes served at church dinners. A proposed definition is too **broad** if it allows things to count as referents that should not count. The following proposal is too broad: An apple pie is an apple-filled dessert. This is too broad, since it includes apple-filled donuts, which should not be counted as apple pies. A definition could be both too narrow and too broad. Consider again the definition the NCIC used when they reported that there were 800,000 child kidnappings in the United States during 2005. Many of us would say that their definition is too broad. We would not include running away from home as a kind of kidnapping. This means that the NCIC's definition left off some

features that we would ordinarily think are necessary for being a child kidnapping. But notice that their definition of a “stereotypical kidnapping” might be too narrow too, since it requires that the child be transported at least 50 miles from home. I doubt that most of us would consider that really necessary for a kidnapping. The NCIC’s definition of “child kidnapping” was, by ordinary standards anyway, both too narrow and too broad.

BOX 2.3.1 PRACTICAL TIP: COUNTER-EXAMPLE METHOD

An important step in constructing or evaluating a proposed definition is looking for **counter-examples**. A counter-example is a case, either a real one or a fictional one, that shows that the definition is either too broad (that is, that it includes things that it should not) or too narrow (that is, that it excludes things that it should not). If you think that a proposed definition is too broad or too narrow, then you need to present a counter-example and argue that it shows that the proposed definition is mistaken. If, on the other hand, you respond to an alleged counter-example to your proposed definition, then you have to either show that it is not a genuine counter-example to the definition or else revise the definition to include/exclude examples of that kind.

2.3.1.4 Definitions in Practical Life How bad is it for a proposed definition to be too broad or too narrow compared with our ordinary standards? While it would be good if everyone had exactly the same necessary and sufficient conditions in mind when we use our words, this is not very likely and probably not needed. Most of us use our words in idiosyncratic ways, ways that are just a little bit different from the way others use them. The things we would count as its referents might not be exactly those that others would count. Just think of the disagreements we would get over whether something is a murder or not. But for practical purposes, the existence of these differences is not important, so long as we are all aware of them. It is less important that we all agree that a definition captures our own idiosyncratic usage than that we all agree on what we will mean by it for the purposes at hand. If we can agree, either just for the sake of the discussion or for good, on what we mean by “child kidnapping,” then this will make it easier to avoid misunderstandings when we start to say what the facts are.

Definitions can report actual usage, but they can also help to standardize it. Both are worthwhile goals.

It is also good for our words to remain flexible to deal with new or unanticipated uses. Sometimes it is better to hold off on being too precise in specifying necessary and sufficient conditions because more research and investigation may be needed before we can decide whether something really is a referent of some word. It took biologists a long time to figure out whether whales fit the definition of a *fish*, and only recently did astronomers decide that Pluto does not fit the definition of *planet*. Making these decisions required finding out more about whales and planets, and thinking hard about what we *want* our words to mean. We usually want our words to leave room to deal with unanticipated uses. Laws, for instance, are usually written so as to leave room for unexpected cases. We rely on judges and lawyers to help us decide how best to apply our words and concepts in cases that the legislators who wrote the laws could not have anticipated. If the words had no flexibility at all, if we insisted on necessary and sufficient conditions from the very start, then we would not know how to describe these new cases. Biologists want their concept of a *species* to be flexible enough to let them describe new phenomena in terms of species. Astronomers decided to adjust the definition of “planet” even though it meant deciding that Pluto was not a planet after all, because they believe that the new definition allowed them to better say the things they wanted to say all along.

BOX 2.3.1.4 MISTAKES TO AVOID: EQUIVOCATING ON THE MEANING OF “WRONG”

It is a mistake to confuse the different meanings of the word “wrong.” We can distinguish at least three dimensions of evaluation that we might have in mind in calling an action wrong. The action might be legally wrong, in the sense that it violates a law or legal statute. The action might be prudentially wrong, in the sense that it has consequences that would be bad for your health or wealth. The action might be morally wrong, in the sense that it violates moral rules or sanctions. These three kinds of “wrongness” are quite different, and an action may be wrong in one sense but not another. Lying to your best friend might be morally wrong, even though it is not legally or prudentially wrong. In defending or evaluating reasons to believe that some proposal is wrong, it is very important to be as clear as one can about what kind of wrongness is at issue.

EXERCISE 2.3.1**1. Short-answer questions:**

- a. What is the difference between a necessary and a sufficient condition?
- b. Could a necessary condition also be a sufficient one? If so, give an example.
- c. Could a sufficient condition also be necessary? If so, explain using the concepts we have discussed in this chapter, and give an example.
- d. Explain how a condition that is sufficient might not also be necessary. Give an example to illustrate your answer.
- e. If you assert that Jones is a millionaire, are you also asserting that Jones meets all of the conditions that are necessary for being a millionaire? Are you asserting that Jones meets all of the conditions that are sufficient? Explain.
- f. What conditions are necessary and sufficient for a conjunction to be true? A disjunction?

2.3.2 Expand

The initial statement of the meaning should be as succinct as possible. It could even be a slogan. Look back at the definition we discussed in Chapter 1 of “knowledge.” We said that according to the traditional philosophical definition, knowledge is justified, true belief. This is about as succinct and slogan-like as possible. But it would be wrong to leave it at that, since there is still plenty of room for misunderstanding. It would be good to expand on the slogan by filling in details and by saying how the concepts used in the slogan are related to one another.

In the case of the definition of *knowledge*, it would be especially helpful to say more about what belief, justification, and truth are, and how they are related. We discussed these issues at great length in Chapter 1. But in a brief definition, it might be enough simply to say something like this:

Knowledge is justified, true belief. To know something, you have to believe that it is the case; if you do not believe it, then you do not know it. But to count as knowledge, your belief also has to be based on enough of the right kind of evidence. That is, it has to be justified. Finally, your belief has to be true. These three necessary conditions are independent

of one another. A true belief might be unjustified; a justified belief might be false; and a person might fail to believe something true even though they have excellent reason to believe it. But all three are needed, and together they are sufficient, for knowledge.

This elaboration of the brief statement of the meaning of *knowledge* fills in the missing details, by making it clearer what is meant by “justification,” and by saying something about how those three key terms are related one to the other.

BOX 2.3.2 PRACTICAL TIP: IDENTIFY CONCEPTUAL RELATIONS

If a definition that you are constructing or evaluating involves several other concepts, it is helpful to explain or consider how those concepts are related. Are they independent of each other? Does one include the others?

2.3.3 Example

It is almost always helpful to provide some examples. But it is not always easy to know which examples to select. Since the goal of providing a definition is to secure common understanding, the example has to be as noncontroversial as possible. That means that it has to be one that everyone involved in the discussion will agree fits the slogan provided. If in constructing a definition you use an example that some people think does not fit, this will only make it harder to secure shared understanding. Remember that the goal is to reach agreement on how to talk about the facts in order to focus on disagreements about the facts themselves.

If the elaboration involves identifying and relating several concepts, as our example above does, then it would be useful to provide examples showing how the concepts are related. For instance, it would be good to supplement the elaboration of our definition of knowledge with an example of someone who has a true belief that is not justified enough for knowledge, or someone who has a false belief that is nonetheless based on pretty good evidence. The examples could be from real life, so long as everyone involved knows the story. But the example could also be fictional. The goal is always to secure mutual agreement on how the language is to be used.

With examples, more detail is almost always better. Suppose one wanted to give an example of a false but justified belief. Here is one:

A child's belief that Santa Claus exists.

We all recognize the idea. But there is not enough in it to make it clear that it really is a case of a justified false belief. It would therefore be better to say the following:

Suppose that little Joan believes that Santa Claus exists, and believes it because her mom and dad have told her that he does, and that he brings presents every year. Joan's belief is justified, since it is based on her parent's testimony and, in general, it is reasonable to believe what your parents say unless you have good reason to doubt them. But her belief is false, since Santa does not exist. So, this is a case of a false but justified belief.

This example is laid out in much more detail, and the detail also makes it really clear how it is an example of what it is supposed to be an example of.

2.3.4 Contrast

Finally, it is usually helpful to contrast the concept being defined with other related ones. It is best to focus on concepts that it is likely that others will confuse for the one you are defining. Since nobody confuses a concept for its opposite, it is usually not very useful to identify the contradictory concept. Indeed, it is not always very clear what the opposite is. What is the opposite of knowledge? Is it false belief, or is unjustified belief the opposite of knowledge? What is the opposite of murder? It is more likely that someone will confuse murder with killing in self-defense or with killing by mistake than with the opposite of murder. It is more likely that someone will confuse child kidnapping with a child's agreeing to run away with a noncustodial parent. Those are the misunderstandings and confusions that definitions are intended to prevent or remedy. In the case of our definition of knowledge, it is better to contrast knowledge with certainty or with consensus opinion, or with mutual agreement. Once again, though, if one chooses as contrasting concepts ones whose meaning is controversial, then the definition will fail to prevent or remedy misunderstanding. The point of working hard to develop a definition is to resolve and avoid misunderstanding, not to generate it.

EXERCISE 2.3

1. Short-answer questions:
 - a. What is the purpose of providing a definition?
 - b. What is a counter-example? What are counter-examples used for?
 - c. When is it acceptable for a definition not to identify necessary and sufficient conditions? Explain why.
 - d. What is it for a definition to be too broad? Give an example.
 - e. What is it for a definition to be too narrow? Give an example.
2. For the following proposed definitions, find a counter-example. Identify whether it shows that the definition is too broad or too narrow.
 - a. Oxygen: a colorless and odorless gas
 - b. Apple pie: a dessert made with apples
 - c. Triangle: a three-sided two-dimensional figure with a 90-degree angle
 - d. Violin: a stringed instrument
 - e. Parent: the father or mother of a human
 - f. Stove: a kitchen appliance used for cooking
3. For the following concepts, compare and contrast the definitions provided in three dictionaries:
 - a. Automobile
 - b. Water
 - c. Tiger
 - d. Honesty
 - e. Knowledge
 - f. To eat
 - g. The tango
4. Using the four-step method, formulate definitions for the following:
 - a. Donut
 - b. Apple
 - c. Honesty
 - d. Regret
 - e. Chair
 - f. Planet

2.4 THINKING CRITICALLY ABOUT FRAMEWORKS

Different disciplines are in part defined by the concepts they use to describe, explain and raise questions about the phenomena they study. Even though geologists and physicists are both interested in earthquakes, they think about earthquakes in different ways. Likewise, even though sociologists and psychologists are interested in family dynamics, they typically employ different concepts for describing, explaining, and raising questions about family life. They employ different **frameworks**, even though they are thinking about the very same phenomena. A framework is simply a set of concepts and methods that define a specific perspective or point of view. Different frameworks allow for different ways of describing, explaining, and raising questions about a phenomenon. Engaging in a discipline requires understanding and being able to think with its framework. Thinking critically while engaging in a discipline requires reflecting on that discipline's framework, on the way its set of concepts is used to describe, organize, and think about the phenomena it studies.

As we saw above in our discussion of necessary and sufficient conditions, it is not always possible to provide a neat definition of a technical definition. Often, this is not even desirable. We often want there to be some flexibility in our concepts, to allow us to respond to new evidence and new discoveries in new ways. This is one reason that it is helpful in providing a definition of a discipline's key concepts to provide examples and contrasts, since providing those can do as much as necessary and sufficient conditions to prevent or remedy misunderstanding.

Let's consider an example: the case of cancer. We can theorize about cancer from many different perspectives. If we think of it from a molecular perspective, then we need to use the concepts of molecular biology to describe cancer. This will include thinking in terms of genes and proteins, and the kinds of processes, structures, and chemical interactions that occur at that level. We can also think of cancer at the cellular level, in terms of the actions and processes that cancerous cells undergo, and how cancerous cells differ from other kinds of cells. Or we can think of cancer at the level of the entire organism, in terms of the animal's internal, systemic responses to cancer and to the operation of the nervous, immune, and reproductive systems. We can also think of it from a sociological level, in terms of how cancer affects family, work, and community relations. In moving from one perspective to another, we are able to describe, explain, and understand aspects of the phenomena that we cannot "see" from the other levels. The other levels lack the vocabulary for describing those aspects.

The feature of frameworks that make them valuable—that they allow us to think about a phenomenon in one clearly defined set of concepts—is also the feature that makes them limiting. There is nothing inherently wrong with this. But it is a mistake to get *stuck in a framework*. This is the mistake of not realizing that there are other perspectives on a given phenomena, problem, or issue. We need to keep in mind that there are always different perspectives on any phenomena, issue, or problem. Indeed, changing perspectives can sometimes lead to solutions to problems that were first identified but could not be solved at a different perspective. If we were not able to think about cancer at the genetic level, our understanding of the causes of cancer would be very much poorer than it is. This is so, even though not everything about its causes can be learned at that level. Sometimes, we need to think about a phenomenon from several different perspectives at once. A doctor who discusses a patient’s cancer only at the cellular level and not also at the sociological or psychological perspectives will not provide a complete treatment. If we think of the problem of urban poverty only from a sociological perspective and not also from the perspective of criminal justice or microeconomics, we are likely to miss or overlook features of the problem that are hard to see from the sociological perspective alone.

BOX 2.4 PRACTICAL TIP: THE RULE OF THREES

When trying to define a problem, it is helpful to think about it from at least three different perspectives. This is especially important when assessing the costs and benefits of a proposed course of action. Deciding how to respond to global warming requires thinking about the problem from economic, fiscal, environmental, employment, and political perspectives, just to name a few.

Sometimes, politicians and interest groups use one framework rather than another when describing a proposed or existing policy in order to influence the public’s attitudes towards that policy. The very same policy is called by one side “drilling for oil” but the other calls it “energy exploration.” In principle, there is nothing wrong with a policy’s being described in different frameworks, since, as we have seen, most policies and problems are multi-dimensional. The search for oil reserves has environmental and economic aspects as well as impacts on employment, on pollution, on the broader economy, on energy

conservation, and on national security, just to name a few. It would be wrong to decide on a policy without having examined it (and its alternatives) from all these sides. There is no privileged perspective. So there is nothing in principle wrong with a policy's being framed in different ways by different politicians or interest groups. But it is a mistake on our part if we fail to realize which framework the policy is being presented to us from within. It is always a mistake to get stuck in a framework.

EXERCISE 2.4

1. For each of the following familiar problems, frame them in the specified way. Your goal is simply to describe the problem using the concepts that are central to the relevant field, not to offer solutions to it.
 - a. The high rates of teenage pregnancy (economic and emotional)
 - b. Religious intolerance (economic and cultural)
 - c. Adolescent drug use (physiological and psychological)
 - d. Online file sharing of music and movies (economic and cultural)

2.5 CLARIFYING BELIEFS AND PROBLEMS

We sometimes have to clarify beliefs and opinions or get clearer about problems that we face. The need to define our beliefs and opinions and our problems is sometimes hidden by the fact that we so often use simple “Yes/No” questions to find out what other people think and we use simple sentences to tell them what we think, even on topics that we all know are very complicated and controversial. Public opinion surveys regularly ask people whether they support this or that government policy or proposed program, and the pollsters are looking for a yes/no answer. “Do you support the war in Iraq?” “Do you support a person's right to own guns?” But we know that our opinions on most topics are very complex. What is more, even when we agree, we might have different reasons for agreeing. On really complex topics—abortion, capital punishment, and teenage drug use—there is room for huge difference in what we believe. But there is just as much room for differences in our reasons even when we believe the same thing. One can support capital punishment for economic reasons, or for political reasons, or out of concern for deterrence, or simply because one

believes that it provides the most appropriate punishment. The same is true in the case of problems. People who agree that some phenomenon is a problem might disagree about what makes it a problem, and in some ways this is even more important than their agreement on its being a problem. We hear talk about “the problem of illegal immigration,” or “the problem of underage drinking,” as if everyone who agrees that those name a problem agree on just what the problem is. In this section, we will study ways to clarify our opinions and beliefs and our problems.

We can use the four-step method to help us to clarify our opinions and the propositions we believe. The basic approach is the very same as with definitions of concepts. It is good to find a neat, brief way to formulate our opinion, and then provide an elaboration of it, focusing on some of the key words and concepts. In some cases, it will be helpful to provide some examples, though in others it will not. In all cases, or at least cases at all interesting, it will be good to mention a few contrasting propositions, ones that it is likely that someone might think is the one you have in mind even though it is not. Here is an example of a use of the four-step method to define a proposition believed:

I believe that lying to friends is always wrong. I think you have a moral duty to answer sincerely when your friends ask you questions. Telling them something that you do not really believe is just wrong. If a friend asks for my opinion on their career choice, I should take their request for my opinion and advice seriously, and tell them what I really think. If a friend asks me for my opinion on which flat screen TV to buy I should be honest with him. I don't mean that you always have to say everything that is on your mind. I think that it is sometimes better to wait until they ask for your opinion before giving it. But if a friend asks you a question it would be morally wrong not to answer it sincerely.

In this example, showing the “SEEC method,” the first sentence states the view in a brief, **slogan**-like way. The next two sentences **elaborate** on it, by making it clear that it is moral wrongness that is at issue, and just what the author has in mind by “lying.” The next two sentences provide a couple of **examples** to help show what the author has in mind. The final three sentences work to **contrast** the author's view with views that are pretty similar but different in important ways. It might still happen, of course, that someone misunderstands the author's view. But if the author has worked hard to make her view as clear as she can, then she will have done her duty as a critical thinker to clarify her view. Notice that in this passage the author does not provide any evidence at all for her opinion. She makes no effort to try

to convince you that she is right that lying to friends is always wrong. Her goal is not to convince you that she is right, but merely to make it clear what her belief is. She is clarifying what it is that she believes, and not offering reasons for why she does or anyone else ought to believe it.

BOX 2.5A PRACTICAL TIP: ASK OPEN-ENDED CLARIFICATION QUESTIONS

When discussing topics with other people, ask them open-ended questions, not questions that allow a “Yes” or “No” answer. This will reduce the risk that superficial agreement will mask interesting and deep differences. Instead of asking:

“Do you think that ...”

“Do you agree that ...”

Ask:

“Why do you think that. ...”

“What do you mean by ...”

“What reasons are there for thinking that ...”

Depending on what the proposition is, you might not need to provide an example. Here is an attempt to define a belief that does not involve an example, and where it is not obvious what an example would be like:

The Montreal Canadiens are the best team in NHL history. I do not mean that they have always won the Stanley Cup, or even always made it into the playoffs. I know that they often struggled. I also do not mean that they have always had the best players, which they plainly have not. I mean that they have the best management, coaches, and fan support system in the entire history of the NHL. They are simply the best-run team ever.

In this example, the author does not provide an example, but she does work hard to contrast what she means by “best team” with several other things that someone might take her to mean. Notice also that she offers those contrasts right after she provides the initial slogan statement of the belief. This is very helpful in this case, since it sets her up

in an elegant way to offer the elaborations of her view. The four-step method should not be thought of as a rigid formula. It is a helpful guide for thinking about what sorts of things to include when trying to clarify your meaning.

We can also use the four-step method to help clarify our view on what makes something a problem. Here is an example:

I agree that illegal immigration is a real problem, but I see no problem with legal immigration. Illegal immigrants are people who live in the United States but do not have legal authorization to be here. That is what makes them illegal. I do not think that legal immigrants are a problem. We should encourage more of them to come to work and live here. Some illegal immigrants are from Mexico, but illegal immigrants can come from all over the world and they pose a problem no matter where they come from. It is not especially Mexican immigrants that I think are a problem. Some people think that illegal immigrants are a problem because they think illegal immigrants are criminals. I am not sure about that. In my view, illegal immigration is a problem because once someone is here illegally we cannot find them to see if we can help them become legal and so help them to make a lasting contribution to our community. This is what makes it a problem.

The author starts by stating his opinion that illegal immigration is a problem, and then in the next few sentences tries to explain what he means by “illegal” and contrasts his view with other closely related ones. Again, he does not offer examples of illegal immigrants though he does in the final sentence mention one thing that he thinks makes being an illegal immigrant a problem. There is a fine line in this passage between clarifying your view and providing reasons to agree with it. Still, it is fairly clear that the author here is trying to state his position clearly as opposed to offering reasons to share it. We can imagine this passage as part of a discussion among people all of whom agree that illegal immigration is a problem. The author would not be trying to convince the other people that it is a problem, since they already agree that it is a problem, but to clarify what in her view makes it a problem.

BOX 2.5B DECIDING WHAT TO DO: CLARIFYING YOUR PROPOSALS

Clarifying and defining are just as important in deciding what to do and in evaluating a proposed course of action. Deciding what to do involves deciding both on an end to achieve and on a means for

achieving it. Both the ends and the means should be clear before one decides what to do. The 4-step definitional method can be useful here too. It is especially important to contrast the proposed end with others with which it might easily be confused. As a handy rule of thumb: if you cannot identify three contrasting ends or means then you have not made the proposed ends or means clear enough.

We have been discussing how to use the four-step method to help us to clearly state our own views and opinions. We can also use it to state another person's views or opinions. This is a good thing to do if we are not sure just what their view is. By writing it out as if it was your own view, and then asking them whether it accurately states their view, you can make sure that you get their opinion. We also need to be able to state another person's view when we wish to raise an objection to it or to her reasons for believing it. In this case, we need to be especially careful that we accurately state her view. If we misstate her view, whether by accident or on purpose, we will have undermined our goal as critical thinkers, which is to try to get at the truth. The four-step method can help us to avoid committing it by forcing us to think hard about how our opponent's views contrast with other closely similar ones.

BOX 2.5B MISTAKES TO AVOID: STRAWMAN FALLACY

It is wrong to distort or misrepresent another person's view or their reasons for their view. When done deliberately in order to score debating points, this is quite rude. But it is just as wrong when it happens through carelessness. While every one has a duty to make her views and reasons clear, you also have a duty to represent her views and reasons as clearly and charitably as you can. When in doubt, ask open-ended clarification questions to increase clarity.

EXERCISE 2.5

1. Using the four-step definition method as a guide, clearly state your views on the following issues. Remember, your goal here is not to give reasons to believe that your view is correct or true, but only to state it in a way that will help others avoid confusing it with other similar views. (If you do not have a fixed opinion, just pretend that you do.)

- a. Whether Sunday is better than Saturday.
 - b. Whether humans descended from other species.
 - c. When abortion should be legal.
 - d. What the country should do to reduce drug use.
 - e. How to deal with the rising costs of college education.
 - f. Whether it is sometimes okay to lie to a friend.
2. Look at the letters to the editor in a newspaper or magazine. Find two or three in which the author is stating his or her view on some issue of interest to you.
- a. Identify the perspective the author is taking on the issue.
 - b. Assess how well the author does at stating that view clearly.
 - c. Propose changes or additions to improve the clarity of the statement.
 - d. Identify two or three contrasting views. Remember, contrasting views are not opposite views; a contrasting view is a similar view that one might easily confuse with the one being stated.

2.6 TECHNICAL DEFINITIONS

Sometimes words get defined in technical ways. This is usually done to avoid misunderstandings and to help resolve debates and disagreements. Sometimes this happens when researchers take a word that has an ordinary use but put it to more rigorous use in their work. This happened with words like “force” and “energy” that now have very specific and pretty clearly defined uses within physics, even though their roots are in our ordinary talk about the universe. Sometimes this can lead to even more misunderstandings.

The debate over whether nicotine is addictive provides a nice case study. In 1964, the U.S. Surgeon General’s Report stated that nicotine is not addictive. In 1988, the Surgeon General’s Report announced that nicotine is addictive. This looks like a pretty substantial and clear cut factual disagreement, as though the U.S. Surgeon General’s office had changed its mind on a scientific topic that should have wide-ranging public policy consequences. But while there was considerable new information in 1988 on the physiological effects of nicotine and smoking tobacco, one relevant factor that changed in the 24 years between the reports was the definition of the word “addiction.” In the 1964 report, the Surgeon General offered the following: “In medical and scientific terminology the practice (smoking) should be labeled habituation to

distinguish it clearly from addiction, since the biological effects of tobacco, like coffee and other caffeine-containing beverages, betel morsel chewing and the like, are not comparable to those produced by morphine, alcohol, barbiturates, and many other potent addicting drugs.”³ This definition of “addiction,” which considered the production of intoxication a necessary condition for an addiction, echoes the definition that was then accepted by the World Health Organization. And it seems right that, given this definition, nicotine is not addictive, since smoking a cigarette does not produce intoxication, at least certainly not like drinking alcohol or using heroin. In 1988, though, the U.S. Surgeon General redefined “addiction,” dropping intoxication as a necessary condition, and held that for a drug to be addictive it is sufficient that it involve highly controlled or compulsive use, produce psychoactive effects, and that its use involve behavior that is reinforced by that use. Under this new definition, nicotine did count as an addictive drug. While there is no doubt that during that 24-year period a good deal more was learned about the science and medicine of drugs and nicotine, it is important when trying to understand that debate to be clear on what “addiction” means. As recently as 1998, the Tobacco Marketing Association published the following: “The definition of addiction is wide and varied. People are addicted to the Internet. Others are addicted to shopping, sex, tea, and coffee. The line I would take is that tobacco isn’t addictive but habit forming.”

The debate over whether nicotine is an addictive substance looks on the surface like a purely scientific and factual one, but there is actually a large terminological element to the debate. All sides in the debate—the cigarette companies, the government regulators, the independent scientists—could reach agreement on all the physiological, psychological and chemical effects, including both the long-term and short-term effects, of nicotine use, and on how those effects are similar to and different from the effects of short- or long-term alcohol or heroin use, and they might still disagree about whether nicotine is addictive, simply because they mean different things by “addictive.”

In itself, this is not unusual or even very bad for researchers or ordinary people to use words in a technical way. But when people do not realize this, there is a risk that the participants will end up talking past each other. This mistake occurs when people in a discussion are using the same words with different meanings and are not aware of this, and

³1964 Surgeon General’s Report: Smoking and Health—Report of the Advisory Committee to the Surgeon General of the Public Health Service; U.S. DHEW 1964 p. 350.

so are not aware what each of them is saying. To avoid this, a good critical thinking strategy is to use the SEEC method to define our own words and those of the people we are in discussion with, in order to ensure that we all know what we mean.

2.7 MEANING IN ADVERTISEMENTS

It is wise to keep an eye out for special definitions in advertisements for goods and services. Given the recent rise in popularity of organic foods, many companies are marketing their products to tap into this popularity. Some are now being advertised as “all natural” or “authentic,” as if this meant the same as “organic.” In fact, the use of the word “organic” is highly regulated by governments around the world. The regulations were put in place mostly to help consumers avoid being tricked by producers and to protect growers whose produce really are organically grown from less scrupulous competitors. But the standards for what counts as “organic” vary from one country to the next. In particular, the United States allows products to be labeled as organic so long as they contain no more than 5 percent nonorganic constituents. One company even tried to market its beers as organic, even though they were made with nonorganically grown hops. Apparently, since most of beer is simply water, and since the hop flavoring is so strong that very little of it is needed, many beers already fit the official definition of “organic.” The government has no choice, if it is to regulate the use of a word, to develop a strict definition including necessary and sufficient conditions. But this has the unintended effect of providing loopholes through which products that ought not to count can slip. In turn, these loopholes can confuse and deceive consumers and penalize producers who are not familiar with the technical definition, even though the original motivation for regulating the use of the word was to prevent consumer confusion and deception. Still, the benefits of regulating the use of that word are probably still higher than its costs.

In the case of the word “organic,” the U.S. government regulates both the meaning of claims used with it and also the truth of those claims. It is against the law to use that word on a product to mean something other than what the government has stipulated it is to mean, and it is against the law to use on a product unless that product meets the relevant standards. But there are many words whose use in advertising the government does not regulate at all. Words used on nutritional supplements and cosmetics are a good example of this. These

advertisements are required, as are all advertisements, not to be deliberately misleading. To comply, the advertisements often avoid strong claims like “will eliminate wrinkles” or “will prevent the common cold” in favor of such weak claims as “will help eliminate the appearance of fine lines” and “can support the sinus and immune system.” In the case of nutritional supplements, these claims are usually accompanied by a tiny footnote that says something like: “This statement has not been evaluated by the Food and Drug Administration. This product is not intended to diagnose, treat, cure, or prevent any disease.” It is hard to know what the intended meaning is of the words “support the sinus system” if it is not to mean that the product can treat, cure, or prevent diseases of the sinus system. This is a case where not only is the truth of those claims not being regulated by the government; neither is the meaning. Here the general advice for thinking about claims in advertising applies: buyer beware.

One strategy for thinking critically about the claims in advertisements is to identify the key words and phrases and look for contrasting ones that might have been used and to ask why they were not used. Why does the facial cream advertisement say that it can help reduce the appearance of fine lines? Why did it say “help reduce” instead of just “reduce.” Why did it mention the appearance of fine lines, instead of just fine lines? Reducing the appearance of fine lines is compatible with the continued existence of those fine lines. Is the ad really only claiming that the product covers up the fine lines? Then why not say that? In this way, the four-step method, which requires us to look for closely contrasting words and concepts, can help us to think critically about claims in advertisements.

Earlier we discussed the way politicians and interest groups pick a specific framework for describing a policy in order to influence the public’s attitude to the policy. Unfortunately, but not very surprisingly, they also sometimes use misleading labels to characterize or name their preferred policy proposals. (Unfortunately, there is no law against misleading advertising of public policies.) After the attack of 9/11, the congress passed a law with huge bipartisan support that made enormous and far-reaching changes to civil rights, the rights of criminal suspects, and the power of the government to investigate and even detain citizens. This law was called the “Patriot Act,” a name that had nothing to do with the substance or rationale for the bill, but which made opposing it rhetorically very difficult. Being opposed to the patriot act is like being opposed to patriotism itself. In a similar way, proposals to cut taxes are sometimes presented as “tax relief,” even

when those benefiting would be multi-millionaires, who can easily afford the taxes they are required to pay. There is, of course, a serious issue about who should pay how much tax. But this debate is stymied if one side is using misleading terms to describe the problem or their proposed solution. As good critical thinkers, we need to stay on guard for this, and the four-step method can help by reminding us to look for contrasting ways to describe a problem or solution.

EXERCISE 2.7

1. Find five ads in your local newspaper.
 - a. Identify the key claims made in the advertisement, and define them.
 - b. Identify three contrasting claims that one might easily confuse for that view.
 - c. Propose changes to improve the clarity of the advertisement.
2. Look for statements by local or national politicians on issues that you care about in newspaper articles, in letters written to local newspapers, or on their websites.
 - a. Identify the framework within which they discuss the issue.
 - b. Proposes changes that would clarify their views.
 - c. Identify two or three contrasting proposal or views.
3. What necessary condition is stated in each of the following?
 - a. If you are going to succeed you need to think hard.
 - b. The audience will love this movie, but only if the action scenes are longer.
 - c. Without more water, this plant is destined to die.
 - d. To make cookie dough, you need sugar, flour, butter, and an egg.
 - e. If you love me, then you will set me free.
4. For the following concepts, find a condition that is necessary but not sufficient, and a set of conditions that are sufficient but not necessary.
 - a. Winning the lottery.
 - b. Being president of the United States.
 - c. Being a doctor.
 - d. Being an illness.
 - e. Being beautiful.

BOX 2.9 CHAPTER SUMMARY

Thinking critically about what to believe or do usually requires reflecting on the meaning of concepts, claims, problems, and proposals. The **Assertion Test** can be used to figure out exactly what a concept or claim means. A useful method for constructing a definition involves providing a **slogan**; **expanding** it by explaining the relations among the slogan's key words; providing an **example** or two; and identifying some **contrasting** concepts, claims, problems, or proposals. Definitions can be evaluated by looking for **counter-examples** showing that the definition is too broad or too narrow. Different disciplines might approach the same phenomena using different conceptual **frameworks**.

2.8 MISTAKES TO AVOID

Mistake of a False Definition It is a mistake for a definition to be too broad (to include things that do not fall under the concept) or too narrow (to exclude things that do fall under the concept defined), or both. A counter-example to a definition is an example that shows that the definition is too narrow or too broad. The four-step method can help us to avoid this mistake by requiring us to look for counter-examples and contrasting concepts.

Equivocating on the Meaning of “wrong” It is a mistake to confuse the different meanings of the word “wrong.” We can distinguish at least three dimensions of evaluation that we might have in mind in calling an action wrong. The action might be legally wrong, in the sense that it violates a law or legal statute. The action might be prudentially wrong, in the sense that it has consequences that would be bad for your health or wealth. The action might be morally wrong, in the sense that it violates moral rules or sanctions. These three kinds of “wrongness” are quite different, and an action may be wrong in one sense but not another. Lying to your best friend might be morally wrong, even though it is not legally or prudentially wrong. In defending or evaluating reasons to believe that some proposal is wrong, it is very important to be as clear as one can about what kind of wrongness is at issue.

Mistake of Talking Past Each Other It is a mistake for participants in a discussion not to recognize that they mean different things by the

key words and phrases they use. This can be recognized and avoided by a careful use of the four-step method.

Mistake of Being Stuck in a Framework It is a mistake to frame or describe a problem in a way that prevents one from seeing certain alternatives or solutions. The four-step method can be helpful in avoiding this mistake, by requiring one to identify closely related but contrasting descriptions of the problem. Sometimes, reframing a problem by describing it in different terms or from a different disciplinary point of view can help resolve it. This is what “thinking outside of the box” requires.

Strawman Mistake It is wrong to distort or misrepresent another person’s views or reasons. Not only does this hurt their feelings (and when done deliberately is exceedingly boorish), but it undermines the goals of critical thinking, which is to seek the truth in another person’s views, even if those views are not our own. The four-step definitional method can help us to avoid this by requiring us when defining another person’s position or belief to identify closely related but subtly different contrasting propositions.

2.9 PRACTICAL STRATEGIES

In this chapter we have seen two practical strategies for helping us to think critically about meaning.

The Assertion Test To tell whether a proposition is among the things a person is asserting or claiming to be true, suppose that it is false and ask whether what the speaker says could still be true. If yes, then that proposition is not among the things asserted; if no, then it is. This test can also be used to tell whether a proposed definition is too broad or too narrow, by considering counter-examples to it.

The Four-Step Method In constructing a definition of a concept, belief, proposal, or problem, it is helpful to formulate it as a slogan, to expand on it by saying more about the key concepts, to offer an example or two, and to provide some contrasting concepts, beliefs, proposals, or problems. The goal of providing a definition is to prevent or remedy misunderstanding. This method can also be used to evaluate definitions.

2.10 FROM THEORY TO PRACTICE: APPLYING WHAT WE HAVE LEARNED

Thinking Critically about Ourselves In Chapter 1, you compiled a list of five or six character traits that you think are essential to being a morally good person, and you wrote a tentative definition of one of them.

- i. Using the four-step method for defining a concept, revise the definition including examples and contrasting concepts.
- ii. Ask three or four friends how they would define the concept. Make sure you ask them open-ended questions to get them to say as much as you can.
- iii. Compare and contrast the definition you developed and the ones your friends provided. What are the differences? Pay close attention to slight differences in word choice, as these often make a huge difference to the definition.

Thinking Critically in the Classroom Thinking critically in a discipline requires knowing how to use its concepts to describe a phenomenon and to frame questions and proposals. In Chapter 1, you compiled a list of five or six of the most fundamental concepts of ideas in your field of study.

- i. Using the SEEC method, try to develop a definition of the concept. Make sure that you include contrasting concepts, which need not be central to your field of study. (You might show your definition to your professor to see whether she considers it too narrow or too broad.)
- ii. List five or six problems or puzzles that your field of study addresses or studies. Using the concepts of your field, define them as clearly as you can, making it clear what it is that makes them a problem or a puzzle.

Thinking Critically at Work Successfully participating in a company or organization requires being able to think about its structure, operations and plans. Employees are regularly faced with problems that need to be solved. They can range from short-term ones that are easily solved (e.g., how to get rid of excess inventory) to longer-term problems that require thinking hard about the organization's goals and

structure (e.g., responding to the lower labor costs from second world competitors). The four-step method can be used to define the problem, thereby helping us to make sure that the solutions we employ are appropriate. It is especially helpful to now and then reframe a problem; even one we have a workable solution for. Looking at a problem from a different direction, or using a different set of concepts for describing it can reveal alternative solutions. At the very least, we will get confirmation that our existing solution is still the best one.

In Chapter 1, you compiled a list of five or six problems that you or your co-workers regularly face at work. Pick one of them and define it using the four-step method, making sure that you make it as clear as you can why it is a problem. Now, try to reframe it by describing it in a different way or from a different perspective (e.g., from a manager's as opposed to a worker's perspective, or from a client as opposed to a manager's perspective). This is probably going to be difficult, since we usually find it difficult to think outside the box.

3

SUFFICIENT REASONS

Critical thinking is reasonable and reflective thinking aimed at deciding what to believe or what to do. As we saw in Chapter 1, part of what makes critical thinking *reasonable* is that it aims at finding or providing reasons for our decisions and beliefs. It is important to have reasons for thinking that our beliefs are true, that we have chosen the proper goals, and that our plans to reach them will be effective. If we don't have reasons, then we will be right only by luck, and relying on luck is not much of a strategy. Of course, we want more than just to have reasons; we want to have good ones. We want our beliefs and decisions to be based on enough of the right kind of evidence. But what does it mean to say that a belief is *based on* or *supported by* certain evidence? How can we tell when we have the right kind of evidence? And how do we know how much evidence is enough? This chapter is all about what it means to have good reasons for your beliefs and plans.

3.1 CRITICAL THINKING AND ARGUMENTS

We can think critically about any subject matter and our thinking can take many different forms. But in every case, thinking critically about

what to believe or do involves providing or considering reasons. As we saw in Chapter 1, critical thinking is also *reflective* thinking and this means in part that thinking critically requires thinking about our reasons *as reasons*. It requires making our reasons explicit and thinking about whether they provide enough of the right kind of evidence for the decisions we have to make. One technique that can help us to be reflective is to put our thinking in the form of an argument. This involves explicitly formulating the reasons we have and then assessing whether they are good reasons. The easiest way to do this is in writing. This chapter is all about analyzing and evaluating written arguments. But before we get to the details, I want to say a bit more about how arguments relate to critical thinking.

Let's start with some examples of thinking about what to believe or what to do. We look out the window to see whether it is sunny; we read a history book to learn what caused the American Civil War; we read the bathroom scale and conclude that our diet is working; we do an experiment and decide that the oxygen is causing the reaction; we conduct a survey and conclude that a majority of the population is probably opposed to capital punishment; we test different recipes to find the best way to make a pie crust; we read the newspaper to see who won last night's game; we make a scaled-down model of our backyard to figure out where it makes most sense to build the new deck; we think about last night's concert to get ready to write a review for the school newspaper; we read the letters to the editor to help us decide what to think about the proposal to build a new bridge. I leave it to you to continue this list, but I trust that it is clear that it can go on for a very long time.

This very partial list gives a sense, I hope, of just how much of our ordinary, daily activity involves thinking about what to believe or what to do. What the examples have in common is that they each involve drawing a conclusion on the basis of reasons or evidence. Our belief that it is sunny out is based on what we saw when we looked out the window; the results of our tests are the evidence we use to decide how to make the perfect piecrust. Having reasons for our decisions is essential to critical thinking. But notice that very few of the examples of thinking about what to believe or do involve reflecting on the reasons we had for our decisions. We probably don't think of what we see when we look out the window as the basis for or as evidence for our belief that it is sunny. Our thinking is not that reflective. We do not formulate the reasons as reasons. I do not mean that this thinking does not involve the use of language. Reading a newspaper to find out who won the game involves using language. I mean that when we decide after reading

the newspaper that our team had better start playing better defense, we probably are not (and usually do not need to be) formulating that conclusion or the reasons we have for it in words. We might be able to provide those reasons if someone asked us to. But formulating them in words is not part of the thinking itself.

BOX 3.1 SUMMARY: ARGUMENTS, PREMISES, AND CONCLUSION

When someone does offer explicitly formulated reasons in support of a claim or a proposal we say that they are offering an **argument** for the claim or proposal. We call the reasons that are being offered the argument's **premises**, and we call the claim or proposal that is being supported, the argument's **conclusion**.

When someone does offer explicitly formulated reasons in support of a claim or a proposal we say that they are offering an **argument** for the claim or proposal. We call the reasons that are being offered the argument's **premises**, and we call the claim or proposal that is being supported, the argument's **conclusion**. If Jones is arguing that we ought to build a second bridge across the river because this is the most cost-effective way to deal with the traffic congestion, then the conclusion of his argument is the proposal that *we should build a second bridge across the river* and the argument's premise is that *building a second bridge is the most cost effective way to solve the traffic congestion*. In this sense, an argument is just a collection of assertions some of which, which we will call the premises, are meant to support one of the others, which we will call the conclusion. Thinking of an argument in this way can help us to see how reasons are related to conclusions and will help us to find some strategies for telling when we have good reasons and we don't.

Sometimes arguments are written down, as in a letter to the editor of the local newspaper where the author offers reasons for thinking that more money should be spent on welfare programs. Books and articles can contain also arguments. Some history books, for instance, contain arguments about the origins of the American Civil War, pulling together evidence of different kinds and from different sources to support the author's conclusion. We can find arguments in the editorials published in the college newspaper. Argument can be spoken out loud, as when a politician offers her reasons for supporting a new bylaw

prohibiting smoking in public places. So long as she is doing more than just expressing her support for the bylaw and is actually saying why she thinks the bylaw is a good one, then she is giving an argument for it. In these cases, someone is explicitly trying to offer reasons in support of a position or belief, trying to persuade others to agree or to at least offer them reasons to agree.

This use of the word “argument” is a bit unusual, since ordinarily by an argument we mean an emotionally heated dispute or disagreement. There is no doubt that discussions about what to believe or what to do can be emotional, and we are often, as we saw in Chapter 1, emotionally attached to our own beliefs, opinions, and traditions. But we also saw that emotional reasons are rarely if ever reasons to believe that something is true or that some practice is good or effective and that emotion can get in the way of thinking critically. The fact that my believing something brings me a good deal of comfort does not by itself show that the belief is true. Our focus in this book is on epistemic reasons, reasons to think that something is true or that our proposals are the best ones. So we can stick with this somewhat unusual sense of the word “argument” because it will help us to think about how good reasons are related to our beliefs and proposals.

Letters to the editor, books, and speeches may contain an author’s reasoning about what to believe or what to do. But they do not always. Sometimes, a person will write a letter to the editor simply to express her opinion on a subject, and won’t include her reasons. Her goal is not to try to convince anyone to agree with her, but merely to make sure that her voice is heard. This is fine, though it would of course be better if we were still undecided or if we are considering a change of mind to know her reasons. Likewise, some speeches are intended to encourage emotions or team spirit and not to provide reasons to believe something. There is no sure-fire way to know when a piece of text or a speech does contain reasoning about what to believe or what to do. We do know that it does if, but only if, the author is offering reasons in support of some conclusion.

The difference between thinking that relies on explicitly stated premises and conclusions and reasoning that does not has nothing to do with what the reasoning is about. Take the case of trying to decide what to believe about the local team’s performance. One could conclude that they are in last place simply by reading the standings published in the paper; or one could reach that same conclusion by reasoning that involved explicit words with premises and conclusions. One could draw the conclusion that the compound is responsible for the chemical reaction from seeing the results of the experiments, or from reading

the published report that lays out the experimental procedures and results. What is common to all cases of reasoning about what to believe or what to do, at least when that reasoning is done well, is that evidence of some sort is being offered in support of some conclusion, and this has nothing to do with whether the premises and conclusions have been explicitly formulated or not. No matter what we are thinking about, if we are trying to decide what to believe or do, then we can put our thinking in the form of explicitly stated reasons and conclusions.

There is a real advantage to putting our thinking in the form of an argument, an advantage that ties arguments very closely to critical thinking. Critical thinking is not just reasonable thinking, it is also *reflective* thinking. Part of what this means is that thinking critically requires paying careful attention to the acceptability and strength of the reasons one is considering. It is true that for a lot of our reasoning we do not reflect very much on whether the evidence is acceptable or strong enough. I look at the bathroom scale and draw conclusions about my weight without reflecting on just what the evidence is that I am relying on, or about under what conditions that evidence is trustworthy. I read the newspaper and draw conclusions about the local politics without thinking very much about what makes that a reliable or reasonable way to form beliefs. But unless we are able to think about these questions our thinking will not be critical thinking. We will be engaged in thinking that we do not fully understand, trusting somewhat blindly, not fully in control of our beliefs and decisions. By reflecting on how reading the newspaper can be a source of evidence, we can take greater control of our beliefs. Instead of letting our beliefs be formed on their own, we will become responsible for them in knowing what their grounds are. Reflecting on the way standing on a scale can give me evidence for my weight; I can become more sensitive to the conditions when that method will provide unacceptable or insufficient evidence. Not only will this help me avoid mistakes about my weight, which is always a good thing, but these lessons may even be transferred to other domains, helping me to develop better and more reliable methods.

Thankfully, we won't always need to formulate our reasons and conclusions in words in order to be thinking critically. But there is almost no harm in doing it periodically and much to be gained from knowing how to do it. The form that our thinking takes and the questions we will need to ask about whether the reasons are acceptable and sufficient may vary from one subject to another. In later chapters we will consider some of this variation in more detail. In this chapter I want to stay focused on what all cases of reasoning about what to believe and what

to do have in common: that they involve offering reasons in support of a conclusion.

EXERCISE 3.1

1. Short-answer questions:
 - a. What is an argument?
 - b. What is the difference between a premise and a conclusion?
 - c. Could an argument have more than one premise? If so, give an example.
 - d. Could an argument have more than one conclusion? If so, give an example.
2. Reread the first paragraph of Section 1.2, in Chapter 1. Try to identify the argument's premises and conclusions.
3. The seventh paragraph of Section 1.2.1 contains an argument. Identify its premises and conclusions.
4. Look at a letter to the editor in a newspaper or magazine. Try to identify its premises and conclusion.

3.2 IDENTIFYING PREMISES AND CONCLUSIONS

We have seen that critical thinking is *reflective* in part because it involves thinking about one's reason as reasons, and considering whether they are **acceptable** and whether they provide **sufficient** support for one's beliefs and decisions. But learning how to distinguish reasons from the conclusions they are meant to support is difficult. It can help to begin by learning how to identify the premises and the conclusion in cases where the reasoning has been made explicit. In this section we will look at how to analyze very simple arguments into premises and conclusions.

Unfortunately, there is no foolproof, sure-fire method for doing this. Authors of editorials may give reasons without saying that they are reasons, or might state their conclusion without saying that it is the conclusion. Sometimes, they even leave their conclusions or reasons unstated and leave it up to the reader to figure out what they are. (Perhaps they do this because they have not themselves reflected enough on the nature of their reasons.) Knowing how to figure out

what the reasons are and what the conclusion is supposed to be is one of those skills that is acquired mostly through careful training and repeated practice. There are lots of examples throughout this chapter to help you practice. But let's look at some examples to bring out some helpful bits of advice.

Consider the following bit of reasoning:

The city should build a second bridge to cross the river, for this is the cheapest solution to the traffic congestion and we should adopt whatever is cheapest.

It might already seem pretty obvious what the conclusion is here. But let's approach the analysis of the text very slowly and methodically, in order to be as reflective as we can. The first thing to keep in mind in analyzing an argument is that reasons and conclusions are always **asserted**. (Well, almost always. Sometimes conclusions or premises are left out altogether. But we'll ignore that for now.) So it is helpful to begin the analysis of an argument by identifying all of the assertions it contains. (It might help to look back at the Assertion Test, discussed in Chapter 2.) In the text above, we can identify three assertions. The first one is expressed in the very first sentence of the text.

1. The city should build a second bridge to cross the river.

The next sentence in the text is a **conjunction**, and you will remember from Chapter 2 that a conjunction makes at least two assertions. So we ought to break this sentence into two. (Our goal at this stage in the analysis is to identify all the assertions that are made. Later we can decide which assertion is the conclusion and which are the premises.)

2. Building a second bridge across the river is the cheapest solution to the traffic congestion.
3. We should adopt whatever solution is the cheapest.

There are several things to notice. First, I left the word "for" out of the second assertion. That word is not really part of the assertion. Rather, it plays a special rhetorical role that I'll return to in a moment. Second, notice that I replaced the word "it" with the phrase "Building a second bridge across the river." I did this because I want each assertion that I identify to be as complete as possible. Even though I know what the word "it" refers to in that sentence, it will be helpful later on if we

replace pronouns like “it,” “he,” “she,” etc. with appropriate names or descriptions. I also left off the word “and” from the third assertion, because it simply indicates the presence of a conjunction, and is not really part of the assertion.

Now that we have analyzed that text into three assertions, we can ask which one is the conclusion. Look back to the word “for,” which I left out of the second assertion. The author used the word “for” to let us know that what comes next is a reason for accepting or agreeing with the point that came first, namely that the city should build a second bridge. The word “for” is usually considered to be a **premise indicator**, because it indicates that something is a premise. More specifically, it tells us that the next thing asserted is going to be a premise. But it also works to tell us that we just got a conclusion. So the presence of that word in that part of the text tells us a lot about the identity of the conclusion and the premises. We know that the first assertion is the conclusion, and the next one is a premise. As it turns out, the third one is a premise too. Eventually we will want to know how those premises are related to each other, but that can wait.

Here is a second text:

A second bridge should not be built, since building one will only encourage more people to drive across the bridge than already do now. What is more, if we build two bridges, then we will end up with traffic congestion troubles on two bridges instead of just one.

We can apply what we just learned to this text. The very first sentence has two parts, separated by the word “since.” We can pretty easily identify two assertions; so let’s separate them.

1. A second bridge should not be built.
2. Building a second bridge will only encourage more people to drive across the river than already do now.

Like the word “for,” the word “since” is a premise indicator. It tells us that we are about to get a premise. When the word “since” occurs, as it does in this case, in the middle of a text, it also tells us that we just got a conclusion. So we know that the first assertion is a conclusion and the second one is a premise. Sometimes the word “since” occurs at the very beginning of an argument. There are examples of this in the chapter exercises. In that case, it tells us that what comes next is a premise, but it does not tell us what the conclusion is. Look now at the final sentence. How many assertions does it contain? (Hint: it is a conditional, and we

discussed what they assert in Chapter 2.) There is only one. Neither the sentence after the “if” (which we call the antecedent), nor the sentence after the “then” (which we call the consequent) is asserted. So we must not separate them into two. This is *really* important, for reasons that we will see later.

3. If we build two bridges, then we will end up with traffic congestion troubles on two bridges instead of just one.

Once again, I left out the words “what is more,” since they are used to indicate the presence of a premise. Using them is like saying: “Here is another reason to believe me.” Now that we have analyzed the text into its assertions, we can ask which the premises are and which the conclusion is. We have in fact already answered that. Once again, the very first assertion is the conclusion, while the next two are premises.

Here is one last example:

It would be too expensive to replace the bridge with a tunnel and tearing the current bridge down will harm all the businesses that rely on cross-border traffic. We should simply build a second span across the river.

Once again, the very first sentence is a conjunction and so should be separated into two assertions. The final sentence is an assertion on its own. So we have three assertions altogether:

1. It would be too expensive to replace the bridge with a tunnel.
2. Tearing the current bridge down will harm all the businesses that rely on cross-border traffic.
3. We should simply build a second span across the river.

But notice that in this case there are no indicator words. There are no words to tell us what the premises are or what the conclusion is. We need to figure it out some other way. The only way left is to use our own judgment. There are different approaches to try. You might start by pretending that you were the author, and asking yourself: If I had written this, which of the assertions would I be trying to convince someone to believe? Which would I want a reader to consider the big take-home message?

Another strategy is to insert an indicator word in between the assertions and see which makes most sense. For this strategy to succeed, it is important to have on hand indicator words that you feel very comfortable with. My own favorite indicator word is “therefore,” which I

know tells me that I just got reasons and am about to get a conclusion. Consider the following:

It would be too expensive to replace the bridge with a tunnel; therefore, we should simply build a second span across the river.

We should simply build a second span across the river; therefore, it would be too expensive to replace the bridge with a tunnel.

When I consider these it seems clear to me that the first one makes much more sense than the second. This suggests that the first assertion on our list is a reason to accept the third assertion. If we tried the same experiment using the second and third assertions we would get the same result. All of this suggests, what might have already seemed a bit obvious, that the third assertion in this text is the conclusion. I call this strategy for identifying premises and conclusions, the *Therefore Test*.

BOX 3.2A SUMMARY: PREMISE INDICATORS

Premise indicators are used to show which assertions are meant to be the premises in an argument. Here is a partial list of premise indicators:

Since, because, for, after all, for the reasons that, given that.

BOX 3.2B SUMMARY: CONCLUSION INDICATORS

Conclusion indicator words are used to show which assertion is meant to be the conclusion in an argument. Here is a partial list of conclusion indicators:

Therefore, so, it follows that, this shows that, in conclusion, this proves that

It is a good idea to find one premise indicator and one conclusion indicator that you feel really comfortable with, in the sense that you know exactly how to use it to indicate to someone when you are about to give them a reason or a conclusion. Having them on hand will help when we are analyzing a text that does not contain them. It will also prove helpful when we try to figure out the relations among multiple premises.

The word “because” is a tricky indicator word. For it is sometimes used in an explanatory assertion, as in the following:

The window broke because the ball hit it.

The reaction ended because the temperature fell too low.

These sentences do not contain arguments. They are explanations. Instead of aiming to provide reasons to believe that something is the case, they aim to say why something happened. And they are either true or false, unlike arguments, which are never true or false. The distinction between an explanatory sentence and a complex sentence containing an argument is a subtle one. But the thing to keep in mind is this: an argument is always trying to give you reasons to believe that something is true, as opposed to claiming that something is true.

**BOX 3.2C PRACTICAL TIP: HOW TO IDENTIFY
PREMISES AND CONCLUSION**

- (i) Identify all the assertions.
 - Analyze conjunctions into their conjuncts.
 - DO NOT analyze conditionals, disjunctions, or sentences with nonasserted nounclauses alone.
- (ii) Look for indicator words.
- (iii) If there are no indicator words, use your own judgment or the *Therefore Test*.

EXERCISE 3.2

1. Short-answer questions:
 - a. What is a premise indicator word?
 - b. What is a conclusion indicator?
2. Compose an argument with two premises using “therefore” as a conclusion indicator.
3. Compose an argument with two premises using “since” as a premise indicator.
4. Compose an argument with two premises using no indicators.

5. In the following texts, identify all of the assertions made and then identify the premises and conclusion.
- a. The infection is getting worse, since the fever is staying high.
 - b. Dinosaurs were mammals and they roamed the earth before humans did. This shows that humans are not the first mammals.
 - c. The leaves are drooping and the petals are falling off. It think that this means that the flower is dying.
 - d. The dress is too short and the wrong color. So, you should not buy it.
 - e. There is still water in the cup. I can see it.
 - f. The team won the game last night. It says so in the newspaper.
 - g. There is no use voting. Politicians do whatever they want, and one vote can never make a difference.
 - h. That plant will die. It never gets any light and it is bone dry.
 - i. Inflation is rising and so are interest rates. I think that a recession is approaching.
 - j. The American Civil War was good for the United States. It clarified the powers of the states in relation to each other and to the federal government, and it highlighted the importance of the constitution's bill of rights.
 - k. When we polled a sample of Americans, 85% of them said they were in favor of some legal limitations on gun ownership. Therefore, most Americans are in favor of gun control.
 - l. According to the computer models, the storm is heading straight for Des Moines. Residents should be warned.
 - m. Either we drive to Florida or we fly. But flying is now really expensive and is also bad for the environment. So, we had better drive.
 - n. Investing in the stock market is like throwing money down the toilet. And we can all agree that that is a bad thing. So investing in the market is bad too.
 - o. The acting was wooden, the scenery was cheap, and the dialogue was empty. That was a terrible movie.
 - p. If you save your money in a bank account you will gain interest. If you gain interest you'll become richer. So, saving money in your bank account can make you richer.
 - q. We have to fire that worker. He is incompetent, always in a bad mood, and he just left early.
6. Using the following sentence as a conclusion, "Tofu is very delicious," compose three arguments with different premise. Use premise

and conclusion indicators. (Note: the arguments do not need to be persuasive, and don't worry about making sure that the premises are true. The argument could be crazy, if you like. So long as they are arguments.)

3.3 DEPENDENT AND INDEPENDENT REASONS

Thinking critically about what to believe or do requires having or finding reasons for the decisions we make. Sometimes, we have more than one reason when we make a decision. When I decided to write this textbook, for instance, I thought that the challenge of writing a book would be fun and rewarding. I thought that I could say things in a helpful and clear way and I was hoping to retire early on the huge profits I would make. (It has been fun and rewarding. One out of three is not bad!) In this case, the reasons were independent of one another, in the sense that each of them would have been a reason to write the book even in the absence of the other two. Looking forward to the challenge would have been a reason to write it even if I was pretty sure that I could not say things clearly and would not be able to retire on the proceeds. In analyzing an argument it is important to figure out whether the premises are independent of one another in this way. There are two strategies we can use to help with this.

3.3.1 The Words Test

The first strategy I call the *Words Test*. To see how it works, consider the following argument:

BOX 3.3.1A PRACTICAL TIP: THE WORDS TEST

There ought not to be more information in the conclusion of an argument than there is in its premises. So, to test whether premises are dependent or independent, look to see whether the conclusion of an argument contains important words that occur only in one premise and important words that occur only in another premise. If so, then those premises are probably dependent.

John should not become a doctor. After all, he really hates to be around sick people, and doctors spend their whole day around sick people.

We can analyze it into three assertions:

1. John should not become a doctor.
2. John really hates to be around sick people.
3. Doctors spend their whole day around sick people.

(Notice that I replaced the pronoun “he” from the second assertion with the name “John.”) We know from what we learned in the previous section that the first assertion is the conclusion. The words “After all” are used here as premise indicators. So we know that the second and third assertions are supposed to provide reasons to accept the conclusion. What we now want to know is whether they are dependent on each other, or whether they provide independent support for that conclusion. We can see that the conclusion contains two important words: “John” and “doctor.” But the word “John” only occurs in the second assertion and the word “doctor” only occurs in the third assertion. This suggests pretty strongly that those two assertions are supposed to be working together—i.e., dependently—to support the conclusion.

Lying behind the Words Test is the following general rule about good reasons: *there should not be more information in the conclusion than there is in the premises*. If an argument’s conclusion does contain more information than there is in the premises, then this means that a logical leap is needed to get from the information in the premises to the information in the conclusion. In other words, the conclusion does not really follow logically from the premises. (We will look more closely at just what that means in a little while, but for now we can stick with a pretty intuitive grasp of that idea.) One way to see whether there is more information in the conclusion is to compare the important words used in formulating it with the important words in the formulation of the premises. That is why the Words Test works pretty well at figuring out whether premises are dependent or not. But it is of only limited use.

3.3.2 The False Premise Test

Here is an argument where the Words Test won’t help:

John is a criminal. He robbed the bank and shot the teller.

We can analyze it into three assertions:

1. John is a criminal.
2. John robbed the bank.
3. John shot the teller.

Using the tests that we learned in the last section, we can see that the first assertion is the conclusion and the other two are premises. Are those premises independent of one another? The Words Test is not going to help us here, since one of the key words in the conclusion, “John,” occurs in both premises and the other one, “criminal,” occurs in neither one. In this kind of case, it is better to use a somewhat more complicated, but much more reliable test.

BOX 3.3.2A PRACTICAL TIP: THE FALSE PREMISE TEST

If a premise would provide some reason to accept the conclusion even if another premise were *false*, then those premises are independent. So, to test whether premises are dependent, suppose that one is false, and ask whether the other one would still provide some support for the conclusion. If it would, then the premises are independent. If not, then they are dependent.

To use this test, we need to suppose for the sake of the test that one of the premises is false. So let’s suppose that John did not rob the bank, but that he did shoot the teller. Just to help us suppose this, let’s further suppose that his friend Michael robbed the bank. Now we ask: would the fact that John shot the teller still be some reason to accept the conclusion that John is a criminal? Clearly it would. Let’s do the test on the other premise. Let’s suppose that John did not shoot the teller, but that he did rob the bank. Would the fact that he robbed the bank be some reason to accept the conclusion that John is a criminal? Again, the answer is clearly yes. What this shows is that those two premises do not depend on each other to be a reason to accept the conclusion. Each is on its own a reason to accept that conclusion.

Let’s use the **False Premise Test** on the argument for the conclusion that John should not go to medical school. Let’s suppose that John *does not* hate being around sick people, but that it is true that doctors spend a lot of time around sick people. Would the fact that doctors spend a lot of time around sick people be a reason to think that John should not go to medical school? No. It is a reason to think that *only if* it is also true that John does not like being around sick people. Try

it the other way. Suppose that it is not true that doctors spend a lot of time around sick people, but that John hates being around sick people. Would that fact about John be a reason for him not to go to medical school? Again, no. It would be a reason for him to avoid becoming a doctor only if doctors spent a lot of time around sick people. Using the false premise test on this argument is a bit trickier, since supposing that that third assertion is false is hard because we know for a fact that it is true. Using the False Premise Test can sometimes require a vivid imagination, whereas the Words Test is a bit more mechanical, requiring only that one look at the words involved. But the Words Test won't always be workable, whereas the False Premise Test will always work.

Let's consider one last case:

Capital punishment should be banned. It is often cruel and cruelty should be banned. Moreover, our justice system sometimes makes mistakes and it would be horrific to execute an innocent person.

We can analyze this into five assertions:

1. Capital punishment should be banned.
2. Capital punishment is often cruel.
3. Cruelty should be banned.
4. Our justice system sometimes makes mistakes.
5. It would be horrific to execute an innocent person.

We know that the first assertion is the conclusion. This leaves us with four premises. Reading them through carefully, we can see that there are two main ideas in the premises. One has to do with whether capital punishment is cruel and the other has to do with whether our justice system might make mistakes. Using the Words Test, we can see that the second and third assertions work together to support the conclusion and that the third and fourth work together to support the conclusion. We can also see, using the false premise test, that these pairs of premises are independent one from the other. Even if the justice system never made mistakes, if it is true that capital punishment is cruel and that cruelty is wrong, then this would be some reason to accept the conclusion. Likewise, even if capital punishment was not cruel, if our system sometimes makes mistakes and if it would be wrong to execute innocent people, then this would be some reason to accept the conclusion. So here we have two independent sets of dependent premises!

BOX 3.3.2B PRACTICAL TIP: DIAGRAMMING ARGUMENTS

It can be helpful to construct a diagram to make the structure of an argument clear.

Give each assertion in the argument a number. Place the number for the conclusion at the bottom of the diagram and the numbers for the premises above. If the premises are dependent, connect with a “+” and then draw an arrow from it to the number of the conclusion. If the premises are independent, connect each of them directly to the conclusion with an arrow.

Dependent Premises



Independent Premises



3.3.3 Circumstantial Reasons

One final point is worth noting here. Sometimes in arguments where there are lots of apparently independent premises, the premises are actually meant to be working together producing a kind of “piling-on” effect. Consider the following argument:

Let’s not go to the movie tonight. I’m tired, and we still have all those dishes to do. Plus, we need to save some money for lunch tomorrow and anyway we can watch a movie on TV. Let’s just stay home.

When we analyze this into assertions we find the following:

1. We should not go to the movie tonight.
2. I am tired.
3. We still have to do the dishes.
4. We need to save money for lunch tomorrow.
5. We can watch a movie on TV.

The first assertion is the conclusion, and the rest are premises. If we used the False Premise Test, we would get the result that each premise constitutes an independent reason to accept the conclusion. This might

have been the author's intent. Maybe he or she thought that each on its own was sufficient reason not to go to the movie. (We will look more closely in a moment at what it means for a premise to be sufficient. But for now we can stick with the rough and ready idea.) But maybe she instead meant that, although none of the reasons is sufficient on its own, when you consider them all together they do constitute a good reason not to go. Maybe she thought that if the dishes did not need doing and if they did not need to save the money for tomorrow's lunch, then the fact that she was tired would not be good reason to miss the movie, but that when all of those considerations are put together, when you pile them all in together, then they do make up a good case. If this were what she is thinking, then there would have to be a missing premise, something like this:

Given all of these considerations, we should stay home.

This "piling on" effect is common in criminal cases, where the prosecution presents a lot of little bits of evidence of guilt, none of which is in itself conclusive, but which, when considered all together, strongly suggest guilt.

John is the robber. He was in the bank at the time of the crime. He owns a weapon of the very same kind as the one used during the crime. He has no alibi for where he was during the crime. And the money from the robbery was found in his building.

Perhaps none of this evidence on its own is conclusive, removing every possible reasonable doubt as to whether John is guilty. But when put together, it does make a pretty strong case. The missing premise here is something like this: If John was at the bank and owns the weapon that was used, and has no alibi, then he is the robber. This makes the argument much stronger, for reasons that we will see later. We will also see later how to identify missing premises. For now, let's just keep in mind that premises are independent of one another if each would be some reason to accept the conclusion even if the others were not.

BOX 3.3.3A DECIDING WHAT TO DO: CLEARLY DISTINGUISH REASONS FOR THE ENDS AND REASONS FOR THE MEANS

Deciding what to do involves deciding on an end to try to achieve and deciding on a means to achieve it. So two kinds of reasons are

important in deciding what to do. For we need to have good reasons both for the ends we are trying to achieve and for the means we are choosing to use. In constructing an argument to represent our reasoning about what to do, we should make sure that we separate out these kinds of reasons. When we analyze someone else's reasons for acting we should also be careful to distinguish the reasons to achieve that end from the reasons to choose those means.

BOX 3.3.3B SECTION SUMMARY

We often have different reasons for our beliefs and decisions. Reasons are independent of one another when each would be a reason even if the other was not. If we explicitly formulate our reasons as an argument with premises and conclusions, then we can ask whether some of the premises work together to support the conclusion. We can use the **Words Test** or the **False Premise Test** to determine whether they are dependent or independent.

EXERCISE 3.3

1. Look back at the argument that you analyzed in Part 6 of Section 3.2. Determine whether the premises are dependent or independent.
2. In the following arguments, identify the premises and conclusions and determine whether the premises are independent or dependent. (Don't worry about whether the premises or conclusion are true.)
 - a. The math class is worth taking because it is easy and the teacher is really nice.
 - b. North Korea is a dangerous country, because it is a dictatorship and all dictatorships are dangerous.
 - c. Jones will probably win the race. He is the fastest skater and the fastest skater usually wins.
 - d. Wind power is the way of the future. It is really inexpensive; it does not pollute; and there will always be wind.
 - e. The plant is dying. The leaves are turning brown, and this is a sign of plant death.
 - f. Raccoons are digging up the grass again, and every time they do this it is because there are grubs. So, the grass has grubs again.
 - g. This camping site already has wood and it has a nice western exposure so we will get a nice sunset. I think this is a good site.

- h.** The new car has higher fuel efficiency and better suspension. We should buy it. Oh, it also comes in five cool colors.
- 3.** For the following conclusions, construct two arguments, one with two dependent premises and one with two independent premises (or two independent sets of dependent premises).
- a.** The team lost the game last night.
 - b.** The liquid contains salt.
 - c.** Harvard University is in Cambridge.
 - d.** The car is out of gas.
 - e.** The rent is due tomorrow.
 - f.** Honesty is the best policy.
- 4.** You probably have reasons in your own life for your beliefs or decisions. Here is a list of topics that are of importance. If you have an opinion, try to formulate it as clearly as you can. (It might help to look back to our discussion in Chapter 2 on how to define a view or position.) Then formulate your reasons, and see whether you have multiple reasons for them. Finally, put them all together in the form of an argument.
- a.** The morality of capital punishment.
 - b.** The morality of abortion.
 - c.** Your decision to go to college.
 - d.** Whether marijuana should be legalized.
 - e.** Whether creationism should be taught in schools.
 - f.** Whether it is sometimes morally OK to lie to your friends.

3.4 SUBARGUMENTS

We have seen that we sometimes have many reasons in support of our beliefs and decisions. Sometimes, we even have reasons for our reasons. In the case of an argument, when an author offers reasons in support of a premise we call that offering a **subargument**. A subargument is simply an argument for a premise. In this section we'll learn how to identify subarguments.

To identify subarguments in texts, we can use the very same methods we studied in the previous two sections. Consider the following text:

The restaurant was not very good. The salad was too salty and the cake was very dry. Worst of all, the service was terrible, for the waiter was slow and the hostess was rude.


Here we find the following assertions:

1. The restaurant was not very good.
2. The salad was too salty.
3. The cake was very dry.
4. The service was terrible.
5. The waiter was slow.
6. The hostess was rude.

Again, it is clear that the first sentence is the conclusion of the argument. Assertions 2, 3, and 4 are independent premises supporting that conclusion. But what about assertions 5 and 6? The presence of the premise indicator word “for” tells us that they are premises. And it tells us that we just got a conclusion. This means that assertions 5 and 6 are meant to provide support for assertion 4, which in turn supports the conclusion. (Assertion 4 is thus both a premise for a conclusion and a conclusion.) If we use the false premise test, we can see that assertions 5 and 6 each provides independent support for 4.

BOX 3.4A PRACTICAL TIP: DIAGRAMMING SUBARGUMENTS

Subarguments can be diagrammed using the “+” and arrow symbols we have already seen. The argument we have been considering would look like this:



Here is a petty reliable rule of thumb. *Unless there is a premise or conclusion indicator that indicates the presence of a subargument, assume that there is not one.* Analyzing texts into component assertions, identifying the main conclusion, and determining whether the premises are dependent or independent is already a lot of hard work. It is pretty

rare, in my experience, for authors of typical arguments—the kinds we find in letters to the editor or even opinion columns in newspapers—to provide a lot of help by using clear indicator words. Perhaps this is because most authors have not thought as much as we have about how to organize their reasons. It seems wrong somehow to work so much harder than the author himself did to reconstruct a nicely structured argument.

BOX 3.4B SECTION SUMMARY

A subargument is an argument for a premise. A subargument can consist of dependent or of independent premises. We can use the same methods for identifying these subpremises and for determining whether they are dependent or independent as we used for arguments.

EXERCISE 3.4

1. In the following texts, identify the conclusions, premises and any subarguments:
 - a. Coke is better than Pepsi. It has more flavor and it is more popular. Survey after survey show that most Americans prefer Coke.
 - b. Joan is definitely pregnant. She missed her period, and the urine test was positive. She told me so herself.
 - c. Slavery is odious. It prevents the slaves from becoming fully autonomous. It also prevents the slave owner from achieving this full humanity, since no one can own slaves and feel good about themselves.
 - d. If we execute Steven for his crimes, then we will set a bad precedent. And setting bad precedents is not good. For it usually leads to unintended consequences down the road. So, we should not execute Steve.
 - e. Building a new bridge is very expensive. The labor will cost a lot, and so will the steel. Moreover, because the trucking industry will have to bypass the city while the construction is going on, the city will lose that source of revenue. We cannot afford these costs. So we should not build a new bridge.
 - f. Downloading music from friends is wrong. It is like stealing from a store, and we all know that stealing is wrong. What's more,

downloading music from friends is effectively robbing from the musician, and we need to support them not rob from them because they are poor and struggling.

- g. If evolutionary theory is correct, then we would expect to see similar bone structures in different species. And we have found this. Dogs and cats have similar bone structures in their arms and paws as birds have in their wings. So, evolutionary theory is correct.
 - h. Abstinence before marriage is a good thing. It promotes more respectful relationships after marriage and before marriage. It helps prevent unwanted pregnancies, which are bad because they make it difficult for mothers and father to achieve their life and career goals.
 - i. Abstinence before marriage is a bad thing. It is important for partners to know whether they are sexually compatible before they marry. What is more, abstinence is just like prohibition on alcohol and that only made the urges even stronger.
 - j. The map says that the park is over on the right. But the map was wrong about the museum's location. It said it was on Elm Street when in fact it is on Green Street. So I am not sure that we should trust the map.
2. One of the exercises at the end of the previous section asked you to identify some of your reasons for your beliefs or decisions. Look back at those, and add some subarguments to them.

3.5 EVALUATING LOGICAL SUPPORT

Critical thinking is reasonable thinking in part because it requires us to have or find reasons for our beliefs and decisions. But we also want our reasons to be good ones. To evaluate whether our reasons are good ones, it is helpful to put them in the form of an argument. Now that we know how to identify premises and conclusions, to distinguish dependent from independent premises, and to find subarguments, we are ready to look at how to evaluate whether reasons are good. We will study a relatively simple and straightforward test to tell when an argument's premises provide the best possible kind of logical support. But first, we need to draw some important distinctions.

Ideally, an argument should have two features. It should have true premises and it should be logically strong, that is, the premises should provide strong logical support for the conclusion. Here is one of the

most important lessons of the entire book: whether the premises in an argument are true has nothing to do with whether they support the conclusion. This may seem pretty counter-intuitive, but it is true and it really is extremely important. To start to see this point, let me illustrate this with some examples that will also help to make clear what we mean by logical support.

Consider the following argument:

(1) Jon Stewart is the prime minister of Canada, and (2) all prime ministers of Canada are Martians, so (3) Jon Stewart is a Martian.

We know that assertion (3) is the conclusion and that assertions (1) and (2) are dependent on one another, as the Words Test or the False Premise Test would show. We also know that all three assertions are false. But this does not yet tell us what degree of logical support they provide the conclusion. Here is the test we can use to evaluate the logical strength of an argument.

BOX 3.5A PRACTICAL TIP. TESTING FOR LOGICAL STRENGTH TEST

To test the logical strength of an argument, suppose that the premises were true. Then ask: how likely is that the conclusion would be true too? The higher the likelihood, the more logical support the premises provide.

The **Logical Strength Test** is really just a variation on the Assertion Test, from Chapter 2. Let's use the test on our argument. First, we suppose that the premises are true. This is a bit hard, since we know for a fact that both are false. But just suppose that there was some other universe or world where they were true, where all the prime ministers of Canada are Martians, and where Jon Stewart is prime minister of Canada. This takes a bit of imagination, but not too much. Now ask yourself: in that world, would Jon Stewart be a Martian? Well, if in that world he is prime minister of Canada and in that world all the prime ministers of Canada are Martians, then he would have to be a Martian too, right? There is, as they say, no escaping that logic. So if the premises of that argument were true, then its conclusion would have to be true too. There is no possible way for its premises to be true and yet for its conclusion to be false. The truth of those premises would guarantee that

the conclusion was true too. This is the very highest degree of logical support. We have a special name for it: an argument is **valid** when it is impossible for the premises to be true and for the conclusion to be false. Validity is the ideal degree of logical support. As we will see later, it is actually not that hard to attain. But before we get to that we need to say more about how validity relates to truth.

Our argument about Jon Stewart shows that an argument can be logically strong and still have false premises and a false conclusion. A valid argument can also have false premises but a true conclusion:

Jon Stewart is a Martian and every Martian hosts a TV show, so Jon Stewart hosts a TV show.

In this argument, the premises are false, though the conclusion is true. If we use the Logical Strength Test, we can see that just as before the argument is valid. If it were true that Jon Stewart was a Martian and that every Martian hosts a TV show, then it would have to also be true that he hosts a TV show. So this argument is valid too. This shows that a valid argument can have false premises and a true conclusion.

A valid argument can also have true premises and a true conclusion:

Jon Stewart is human and all humans have parents, so Jon Stewart has parents.

There is only one possibility that validity forbids: no valid argument can have true premises and a false conclusion. This simply follows from the very definition of validity. Validity means that it is not possible for the premises to be true and the conclusion false. So, if you are evaluating an argument and you know that the premises are in fact true and the conclusion is false, then you can safely conclude that the argument is not valid:

Jon Stewart is a man, and all men are humans, so Jon Stewart is prime minister of Canada.

This is a silly argument, one that no one would take seriously. But it does illustrate the point that an argument with true premises and a false conclusion cannot be valid.

Keeping these points about validity in mind is more than a little tricky. I suggest that you find examples of each combination we discussed that you feel totally comfortable with, and which you can compare to other arguments in order to see whether those other

arguments are valid or not. The examples I have given might do the trick. But if they do not, then you should come up with your own.

**BOX 3.5B DECIDING WHAT TO DO:
EVALUATING REASONS**

Evaluating reasoning about what to do requires independently evaluating both the reasons for pursuing the end in question and the reasons for adopting the proposed means. Reasons to pursue some goal or end are reasons to believe that the end or goal is a good one. Reasons to adopt some means to achieving that end are reasons to believe that those means will be effective in bringing about that end. We can make these reasons explicit in the form of an argument.

Here is how Jones might reason about becoming a lawyer:

I want a career that will bring me money and responsibility and also be fulfilling. So, I should become a lawyer or a doctor. But I do not like to see blood, so I should not become a doctor. So, I should become a lawyer.

Here is how Jones might reason about the appropriate means to achieve that objective:

The best way to become a lawyer is to get accepted at the best law school I can afford. If I study hard for the LSATs, do lots of extracurricular activities, and volunteer on the weekends, I will have a good chance of getting into the University's law school. So, I should do that.

These arguments can be assessed as valid or not.

EXERCISE 3.5

1. Short-answer questions:

- a. What does the word "valid" mean?
- b. If an argument is valid, must its premises be true? Using the concepts discussed in this chapter, explain your answer. Give an example.
- c. If an argument has false premises, must it be invalid? If not, give an example.

2. Using the following proposition as a conclusion, “Tofu is delicious,” construct three argument:
 - a. One that is valid and has two false premises
 - b. One that is invalid with two true premises
3. Using the Logical Strength Test, assess whether each of these arguments is valid:
 - a. Cats are warm-blooded and warm-blooded animals are mammals, so cats are mammals.
 - b. The table is blue, so it is colored.
 - c. The War of Independence was a revolution, and revolutions are morally wrong, so the War of Independence was morally wrong.
 - d. If a plant dries out it will die. This plant is all dried out. So it will die.
 - e. I should make dinner. It is my turn and my wife and I take turns.
 - f. The cat is asleep. Cats always dream when they are asleep, so he is dreaming now.
 - g. $2 + 2 = 4$ and $4 + 4 = 8$, so $2 + 2 + 4 = 8$.
 - h. Lying to someone is like robbing them of the truth, and robbing is wrong, so lying is wrong too.
 - i. The movie was terrible. It was too long and the theatre was way too overcrowded.
 - j. Running helps to build cardiovascular strength and can extend you life. Anything that has these effects is good for you, so running is good for you.
4. For each of the following, find a conclusion that follows validly and one that does not.
 - a. The table is made of wood and wood always dries out.
 - b. Cats are warm-blooded, and warm-blooded animals eat meat.
 - c. Jones is a bachelor.
 - d. Frank murdered Henry.

3.6 MISSING PREMISES

If an argument is not valid, then it is not a good argument. Of course, as we know, an argument that is not valid might still have a true conclusion. So, even if we have determined that an argument for some conclusion is not valid, we are not yet in a position to assess the truth

of the conclusion. But sometimes, an invalid argument can easily be transformed into a valid one, simply by adding a premise that the author might have had in mind, but left out. Here are some examples:

Miranda really wants her plants to thrive, so it will probably rain today.

It has not rained in many days; so it is will probably rain today.

The bridge is too expensive. We should not build it.

The Logical Strength Test makes clear that these arguments are not valid. It is not that hard to imagine a world where even though Miranda is praying hard for rain the drought continues, or where even though it has not rained in many days it still won't rain today, or where even though the weather report calls for rain it will not rain. So it is possible for the premises to be true and the conclusion to be false.

BOX 3.6A PRACTICAL TIP: BE CHARITABLE

When reconstructing someone's reasons, it is best to try to turn it into a valid argument. This may require adding a missing premise. The reason to reconstruct arguments so that they are valid is that this focuses the discussion on whether the premises (including the ones that you added) are true, as opposed to the question of whether the premises are sufficient. As we know, it is easy to add a premise to make an argument valid. But this alone will not make an argument a good one: its premises must be true or acceptable as well.

There is a trick for turning any invalid argument into a valid one. Simply add to the argument as a new premise a conditional whose antecedent is the existing premise and whose consequent is the conclusion. Using our previous examples, we get the following valid arguments:

Miranda really wants her plants to thrive, and *if she really wants her plants to thrive, then it will rain today*, so it will rain today.

It has not rained in many days, and *if it has not rained in many days, then it will rain today*, so it will rain today.

The bridge is too expensive, and *if it is too expensive, then we should not build it*; so, we should not build the bridge.

These arguments are now valid, as an application of the Logical Strength Test will confirm. In evaluating them, there is now no question of whether the premises logically support the conclusion: they provide the strongest possible kind of support. But remember that an argument is good only if both its premises support the conclusion and its premises are true. Transforming an invalid argument into a valid one won't necessarily make it into a good argument. For the premises might still be false. Indeed, it might be that only a false missing premise would make it valid. Such an argument is obviously a bad one. So the question of the truth of the premise remains. And now that there are two premises, we can ask of each premise whether it is true. And when we ask, of the new second premise, whether it is true, we are simply asking, in a slightly different way, the same question we asked when we considered how much logical support that original premise on its own provided for the conclusion. Still, by focusing our attention on the truth of the premises, rather than on the question of support, we can get clearer on the factual questions at issue, and this is (I think) always a good thing.

EXERCISE 3.6

1. Using the four-step definition method from Chapter 2, construct a definition of validity that would help someone who had never studied critical thinking understand it.
2. Using the concepts you have learned in this chapter, explain why it is better for an argument to be valid than for it not to be valid.
3. In Chapter 1, we saw that emotional reasons are not epistemic ones. Using the concepts from this chapter, explain why this is so.
4. If an argument has a false premise, might it still be valid? Using the concepts from this chapter, explain your answer, and give an example.
5. Using the four-step definition method from Chapter 2, define the following:
 - a. Premise indicator
 - b. Independent premise
 - c. Subargument

6. The following arguments are not valid. Using the Logical Strength Test, and the concepts we have learned in this chapter, explain which of the pairs is strongest.
- a.
 - i. The glass is full of water. I can see it with my own eyes.
 - ii. The glass is full of water. Joan told me so.
 - b.
 - i. The car is really low on gas. We have to stop.
 - ii. The car has a flat tire. We have to stop.
 - c.
 - i. The sun will rise tomorrow. After all, it has risen every day for the past million years.
 - ii. The sun will rise tomorrow. For the Earth continues to spin on its axis.
7. For each of the following conclusions, construct three arguments providing different degrees of support for that conclusion.
- a. The LA Lakers are the best hockey team in the NHL.
 - b. The dress is too short.
 - c. The economy is losing steam.
 - d. That is a bird.

BOX 3.6 CHAPTER SUMMARY

Thinking critically about what to believe and what to do requires having reasons. We can make these reasons explicit in the form of an argument, with the reasons as premises and the belief or action as the conclusion. A valid argument is one where if the premises are true the conclusion has to be true too. We may have independent reasons to believe or do something. We may also have reasons to believe these reasons, which we may make explicit in a subargument. Sometimes, reasons for believing or doing something are left unsaid or implicit. Indicator words are useful in constructing and analyzing arguments.

3.7 PRACTICAL STRATEGIES

We discussed two strategies for deciding whether an argument's premises are dependent or independent:

Words Test: If the conclusion of an argument contains important words that occur only in one premise and important words that occur only in another premise, then those premises are probably dependent.

False Premise Test: if a premise would provide some reason to accept the conclusion even if another premise were *false*, then those premises are independent.

We discussed a test for deciding how much logical support an argument's premises provide:

Logical Strength Test: To test the logical strength of an argument, suppose that the premises were true. Then ask: how likely is that the conclusion would be true too? The higher the likelihood, the more logical support the premises provide.

3.8 FROM THEORY TO PRACTICE: APPLYING WHAT WE HAVE LEARNED

Thinking Critically about Ourselves In Chapter 1, you identified some character traits that you think are essential to being a morally good person. You picked one of them and wrote out some reasons for thinking that it really is essential. In Chapter 2, you then provided a definition of the trait. Now that we have studied the nature of arguments and the structure of reasons, do the following:

- a. Construct two arguments for the conclusion that that character trait is really essential to being a morally good person. Make sure that you rely on the definition you developed in Chapter 2. Make sure that the arguments are valid.
- b. Construct an argument for the view that it is not really essential. That is, the conclusion of this argument should be that one can be a morally good person even though you lack that trait. Make sure that the argument is valid. Try to make this argument as strong as you can, by using premises that are reasonable.

Thinking Critically in the Classroom In Chapter 1, you developed a list of the five or six most important concepts in your chosen field of study. (If you do not have a chosen field of study yet, then just pick your favorite course.) Look for arguments in the texts you use in that course. Find five arguments. Analyze them into their component asser-

tions. Identify the premises and conclusions. Identify any sub-arguments. Rewrite them in such a way that their logical structure is perfectly clear. Make sure that the argument is valid.

Thinking Critically at Work Studies show that employers value an ability to think critically more than just about any other trait in an employee. They want their workers to be able to think critically about both day-to-day problems as well as about broader organizational performance and plans. In Chapter 1, you listed several tasks that you do at work. Pick one of them, and do the following.

- a. State what its goal is.
- b. Identify three reasons for thinking that that goal is valuable for your organization, and formulate those reasons into a valid argument.
- c. Identify reasons for thinking that the task you chose will in fact succeed at attaining that goal, or is an essential part of what it will take to attain that goal. Formulate those reasons into a valid argument. (These are challenging; do your best, and do not worry too much about accuracy here. The goal of the exercise is to get you to think critically about what you did at work.)

4

ACCEPTABLE REASONS

Critical thinking is reasonable and reflective thinking aimed at deciding what to believe and what to do. It is reasonable in part because it requires us to have or to find reasons for our beliefs and decisions—reasons to think that our beliefs are true or that our decisions are the right ones. Critical thinking is reflective in part because it requires us to think about whether our reasons are good enough ones, and this means that to think critically we need to think about our reasons as reasons. We saw in Chapter 1 that good reasons are ones that are both **sufficient** and **acceptable**. Chapter 3 was all about what it means for reasons to be sufficient. We saw that whether our reasons are sufficient depends on whether they logically support our beliefs and we saw that, somewhat surprisingly, this has nothing much to do with whether they are true or acceptable. Whether our reasons are sufficient depends only on whether, if they were true, that would make it likely that our beliefs would be true too. But obviously we want more than just to have reasons that logically support our beliefs; we also want those reasons to be acceptable or true. This chapter is all about what it means for reasons to be acceptable.

Sometimes, the reasons we have for our beliefs are themselves supported by other reasons. As we saw in Chapter 3, arguments sometimes have subarguments, where one of the premises in the argument is also a conclusion from other premises. Jones believes that the city should build a new bridge, because he believes that this is the best way to solve the traffic problems, and his belief that this is the best way to solve the traffic problems is a conclusion he based on his study of the different ways that other cities have tried and failed to solve their traffic problems. His belief that the city should build a new bridge is based on reasons that are themselves based on reasons, and this creates a sort of structure or web of reasons. I think that this kind of case is pretty common. In cases like this, the question whether those reasons are acceptable is in part the question whether that subargument is any good, and this in turn is simply a question of whether the premises in that subargument are both sufficient and acceptable. As you can imagine, this can go on for a while, with premises supporting premises supporting premises and so on.

But sometimes our reasons to believe something come directly from a source of evidence. This happens when we rely on observation, measurement, or testimony for our evidence. We figure out whether it is raining by looking out the window; we determine our weight by reading the bathroom scales; we predict the future direction of the local economy by reading the newspaper. In these kinds of cases, we rely on a source of evidence to supply us with reasons for belief and action. Observation, measurement, and testimony are among our most valuable sources of evidence. Indeed, if we did not have such sources of evidence, we probably would not have any justified belief (let alone knowledge) at all. It is a traditional if not quite a universal view among philosophers that all of our knowledge must ultimately rest in some way or another on evidence supplied by such sources. Whatever the merits of this view might be, what is important for our purposes is that in order to decide whether a decision that relies on such evidence is good, we need to know how to decide when the evidence provided by a source of evidence is acceptable.

In this chapter we will focus on three sources of evidence: observation, measurement, and testimony. These sources differ from one another in important ways, and this means that the questions we need to ask in assessing the evidence from perception are not all the same as those we need to ask in assessing the evidence from measurement or testimony and some of the terminology we need to frame these questions varies too. We will study these questions and cover this

terminology. But before we look at those different sources, it will be helpful to begin at a somewhat more abstract level. For no matter what source of evidence we consider, the fundamental question to ask is the same: Is the source of evidence **reliable** in this particular case? Let's take a closer look at what this questions means, and how we can answer it in a particular case.

4.1 RELIABLE EVIDENCE

4.1.1 Reliability

A source of evidence is reliable when it provides accurate or truthful evidence more often than not. We can illustrate the idea of a reliable source of evidence with a somewhat silly example. Suppose that we want to know whether it is currently raining in Washington, D.C. Here are two sources of evidence. We might use the coin-flip method which tells us the following: flip a coin and if the coin lands heads, believe that it is raining; if the coin lands tails, believe that it is not raining. We might instead use the phone-call method: phone our friends who live in D.C. and ask them whether it is raining. If they say that it is, then believe that it is raining; if they say that it is not, then believe that it is not raining. It is pretty obvious that the phone-call method is better than the coin-flip method. But why is it better? Why is a friend who lives in D.C. a better source of evidence than the result of randomly flipping a coin? Part of the answer, surely, is that the phone-call method is far more likely than the coin-flip method to provide us with accurate evidence about the weather in Washington, D.C. The coin-flip method might give us the right answer on this particular case. But this would only be by sheer luck. And the chance that it would give us the right answer next time is not very high. In fact, the chance that it would give us the right answer is probably no better than 50%. By contrast, the chance that the phone-call method would give us the right answer is much, much higher, maybe even as high as 95%. So part, anyway, of what makes the phone-call method a better source of evidence than the coin-flip method is that the phone call method is far more likely to give us the right answer. This is part of what makes the phone-call method a better source of evidence than the coin-flip method, at least concerning the weather in Washington, D.C. This silly example nicely illustrates how the acceptability of evidence depends on whether it

comes from a reliable source. Because the phone-call method is more likely to give us the right answer means that it is more reliable than the coin-flip method. The more likely it is that a source of evidence will yield us the right answer, the more often it yields the truth, the more reliable it is.

There are three related points to keep in mind when thinking about the reliability of a source of evidence. First, the reliability of a source of evidence is a **matter of degree**. Some sources of evidence are more reliable than others. It is hard to imagine a source that could be 100% reliable, since every source we know about can malfunction. So we should not demand that level of reliability. But it is plausible to think that there is a minimum level of reliability needed for the evidence provided by a source to be acceptable. And it seems right that the coin-flip method will never yield acceptable evidence simply because it is not nearly reliable enough.

The second point is that the degree of reliability of a source of evidence depends on whether it is operating in **optimal conditions**. My friend in D.C. might be a better judge of the weather in the morning on his way to work than in the afternoon when he is sitting in his office cubicle. This means that as a source of evidence on the D.C. weather he is more reliable in the morning than in the afternoon. The bathroom scale is reliable only when it is on a flat and level surface. Our eyesight is reliable only when the lighting is right. Optimality conditions can vary from one source of evidence to another. In asking whether a source of evidence is reliable we need to keep in mind whether the conditions are optimal. This means that we cannot judge whether a source of evidence is reliable unless we know whether it is operating in optimal conditions, and this means knowing what those conditions are. The more we understand about how our sources of evidence work to provide us with evidence, the more reflective we can be in our thinking.

The third point to keep in mind is that reliability is **topic relative**. A source of evidence may be reliable on some topics but not others. My D.C. friend might be a reliable source of evidence on the weather in D.C. but not on which wines to have with fresh fish. Maybe he can tell when it is raining, but not whether it is better to have a German or a French white wine with grilled ocean salmon. The bathroom scale is a good source of evidence on my weight, but not on my mood or cholesterol levels. This means that when we judge whether to accept some evidence we need to ask whether the source is an appropriate source on this topic. Just because the source is reliable for one topic does not mean that it is reliable for others.

BOX 4.1.1 SUMMARY: RELIABILITY

A source of evidence is reliable just in case it provides accurate information most of the time. The reliability of a source of evidence is always a matter of degree, depends on optimal conditions, and is topic relative.

4.1.2 Undermining and Overriding Evidence

These points about reliability raise two important further questions. First, how reliable does a source have to be for the evidence that it provides to be acceptable? And, second, do we have to know whether the source is reliable in order for the evidence that it provides to be acceptable? These are difficult questions and philosophers and experts in critical thinking have answered them in different ways. Some think that evidence is acceptable only if we know that its source is extremely highly reliable. Others think that it is enough for the source to in fact be pretty reliable and that we do not need to know whether it is for the evidence it provides to be acceptable. So long as it is in fact reliable, this is enough for the evidence it provides to be acceptable (setting aside the question of consistency, which we will consider in a moment.) These are complex topics. Thankfully, we do not need to answer these questions here. Instead, we can make do with a negative rule of thumb. The evidence provided by a source of evidence is unacceptable if we have good reason to believe that it is **undermined** or **overridden** by other evidence we have. Let's look at both these ideas.

Evidence is undermined if we have good reason to think that its source is not reliable. Suppose that I am using the phone-call method to decide whether it is raining in Washington, D.C., but that I then discover that my friend in D.C. has a twisted sense of humor and takes special pleasure in tricking people. This means that I have good reason to think the information he is giving me might not be accurate. It does not necessarily mean that he has been lying to me, or that he will lie to me in the future. He might, for all that I know, take his weather reporting duties extremely seriously and would never lie about the weather. Nor does it mean that the evidence he is giving me is false or inaccurate. Still, because I have reason to doubt his reliability, the evidence he is providing me is not really acceptable.

Consider another case. Suppose that my kids have been playing with the bathroom scale again and that I know that the last time they did this, the scale broke and gave crazy readings. I now have reason to think

that they have once again broken the internal mechanisms. In that case, I should not trust what it tells me when I step on it. Of course, the reading I get from it might be accurate, but still I should withhold judgment until I can make sure that it is not broken. (Doing this requires calibrating the scale with one I know—or at least have no good reason to doubt—is accurate.) In both these cases, the evidence provided by the source is undermined by evidence I have that the source is not reliable. The weather report from my friend, though usually reliable, is undermined by my reasons for thinking that he is tricking me; the reading from the bathroom scale, though usually reliable, is undermined by my doubts about the internal mechanisms.

As we noted, evidence that is undermined might nonetheless be accurate. The bathroom scale might be perfectly fine even though I have reason to think that it is unreliable. (Indeed, the scale might in fact be reliable, even though I have good reason to think that it is not reliable. We will look at more examples of this in a while.) It is one thing to have reason to doubt the reliability of a source of evidence and quite another to have reason to conclude that it is supplying false evidence. (This is a subtle but extremely important point. We will return to it more than once.) But when we do have reason to think that the evidence supplied by some source is false, we say that that evidence is overridden. This happens when we have conflicting or inconsistent evidence; when we have pieces of evidence that point in different directions. Two pieces of evidence are consistent with one another if they can both be true or accurate. If my D.C. informant tells me that it is sunny but the weather station reports that it is actually raining, then I have inconsistent pieces of evidence. One or maybe both bits of evidence must be mistaken. Obviously, in a situation where we have conflicting evidence like this, where we have evidence points in different directions, we have to be extremely careful. It will not always be obvious which source of evidence is at fault. We can formulate some general rules of thumb. A piece of evidence is overridden if:

- it conflicts with evidence from a known reliable source; or
- it conflicts with expert opinion; or
- it conflicts with what we already have good reason to believe.

In cases of conflict, we need to make a judgment. Do we reject the new evidence; do we reject the evidence it conflicts with; or do we reject both? It will not always be obvious which piece of evidence is most accurate. Maybe the source we thought was reliable has made a mistake.

Maybe the expert opinion is wrong in this case. Maybe the proper response to the new evidence is to make a relatively large revision to our standing beliefs. If we want to avoid making a mistake, we should withhold judgment altogether, and to collect more evidence before we make a judgment. We need to decide which source of evidence is most reliable. Until we make that decision, the best or at least the most prudent course is to withhold judgment as long as possible.

BOX 4.1.2A MISTAKES TO AVOID: RELYING ON IGNORANCE

It is a mistake to accept some evidence just because one does not know of any undermining or overriding evidence. It would be a mistake to treat evidence as acceptable or true just because one has not found any conflicting evidence or any evidence that its source was unreliable. One has an obligation, before deciding what to believe or what to do, to look for overriding and undermining evidence. Fulfilling this obligation is part of being a good critical thinker, because it is part of developing a fully rounded case for the belief or plan in question.

We have seen that evidence is acceptable if it comes from a reliable source and it is neither (ii) undermined nor (iii) overridden by other evidence we have. We usually do not need to know that a source is reliable in order to be justified in relying on it. But if we have reason to think the source is not reliable or to think that the evidence is inaccurate, then we are not justified in relying on it. These general points about when we can trust evidence from a source apply to every source of evidence. They are general enough that we can keep them in mind whether we are relying on observation, testimony or measurement. This will help us as we move forward. Still, there are important differences between these three sources of evidence, and seeing them will help us identify some additional questions to ask. We are now ready to consider them.

BOX 4.1.2 SUMMARY: UNDERMINING AND OVERRIDING EVIDENCE

Evidence is **undermined** if there is good reason to think that its source is not reliable, either in general or in this particular case.

Evidence is **overridden** if it conflicts with evidence from a known reliable source, if it conflicts with expert opinion, or if it conflicts with what we already have good reason to believe.

EXERCISE 4.1

1. Short-answer questions:

- a. What does it mean for a source of evidence to be reliable?
- b. Why is reliability a matter of degree?
- c. Explain why reliability is topic relative? Could there be a source of evidence whose reliability is not topic relative? Explain.
- d. What is the difference between overriding and undermining evidence?
- e. Construct an example of a case in which some evidence is overridden but not undermined.
- f. Could a source that is highly reliable nonetheless provide false evidence? Describe an example other than the ones discussed in the text.
- g. Suppose that you had evidence that undermined the evidence provided by some source of evidence, S. Could it still be that S is highly reliable? Explain, and use examples to illustrate.
- h. Suppose that the evidence provided from one source always conflicted with the evidence provided by another source. Should we continue to trust those sources? Which one should we doubt?
- i. Some people think that fortune telling is a good source of evidence. What do you think? Why?
- j. Suppose that we wanted to determine whether perceptual observation is a reliable source of evidence about the colors of medium-sized objects. How could we do this?

4.2 OBSERVATION

The seventeenth-century English philosopher John Locke claimed that without perceptual observation we would have no ideas or thoughts and so no knowledge at all. It is very difficult to disagree with this claim. From the moment we wake up in the morning, we rely on our observations of our surroundings to get around—to find out where we are and what we have to do to get our breakfast. We know that no one sense

organ is necessary—blind people still have beliefs and knowledge, as do deaf people and people like Helen Keller who lack several senses. But it is hard to imagine how a person who had no sense organs at all could possibly have any knowledge of anything at all. Perceptual observation certainly seems essential to knowledge, or at least to human knowledge. At the same time, we know that observation is not infallible—it can and sometimes does make mistakes. Sometimes, things are not quite as they seem to be. So while we have little choice but to rely on our observations, we need to do so reflectively. In this section, we'll study the conditions under which we are justified in relying on perceptual observation.

By perceptual observation we can include the ordinary five senses—taste, touch, smell, hearing, and sight. But we can also include our capacity to tell such things as when we are hungry or thirsty and to tell the relative position of our body parts, such as where our arms are in relation to each other—a capacity called “proprioception” by philosophers and psychologists. Each of these sources of evidence about the world and ourselves is reliable, but only under certain conditions. Sight, for example, only works properly when the external conditions are right. The light has to be bright enough, but not too bright. Changing the color of the light can affect the visual appearance of things. There are also internal conditions that must be right. Vision is not reliable after the optometrist has put dilating drops in—everything looks fuzzy and shadowy. The internal and external optimality conditions are already pretty familiar to us, and we do not need to go into a lot of detail about them here.

Perceptual observation is reliable on some topics but not others. We can often tell by looking what colors things are and where a thing is in relation to other things. We can tell whether the toaster is on the counter or in the cupboard. We can tell by hearing whether the radio is on. We can tell by smell whether the stew is burning. We also know that different senses are reliable on different topics. We can tell by looking, but not by smelling, whether the TV is still on. We can tell by smelling, but not by hearing, whether the milk has gone sour. We can tell by touch, but not by sight, whether the water is cool enough for swimming. It is obviously important to make sure that we are using the proper sense for a given topic. It is also familiar that there are lots of things we cannot tell with any of our sense organs. We cannot tell by observation whether a person has AIDS—we need to run complex tests for this. Nor can we tell by observation alone whether the economy is improving—we need to make some complex measurements for this. Running the tests and making these observations would be impossible

without observation—but the evidence they yield is not observational evidence.

Observation improves as we mature and with training. Anyone who has tried to teach little kids to swing a baseball bat knows that it can be frustrating. It seems to take little kids a long time to learn how to time the bat's swing—something that seems so easy for adults. It is as if the kids cannot even see the ball. Recent studies on the development of the visual system seem to suggest that this is exactly what is going on! The capacity to tell how quickly things are moving requires a relatively advanced level of brain development. On reflection, this is not that surprising. Our perceptual systems are after all just part of our body and we know that our bodies mature and change. In fact, it would be surprising if our perceptual systems did not become more reliable as we grew up. We also know that we can increase the reliability of our perceptual systems with training. Trained musicians can hear rhythms and melodic progressions and patterns in musical performances that others cannot hear. Skilled gardeners can tell by looking whether plants need watering or fertilizer. Doctors learn to identify various skin conditions just by sight. Experts on wine can taste things in wine that most of the rest of us cannot. These improvements are not just the result of maturation—they result from training and practice.

Perceptual capacities are subject to illusions. This is especially familiar in the case of vision. Some visual illusions are **optical**—that means that their explanation has to do with the way light works. For instance, a straight stick in a glass of water looks bent because the light reflected off the part of the stick in the water is slowed down as it travels through the water causing it to change directions slightly, causing the illusion. But other visual illusions are **cognitive**—they have to do with the way our visual system is structured or the way it works. The illusion that parallel railroad tracks meet has to do with the distance between our eyes. The same is true of the Meyer-Lyer illusion, we saw in Chapter 1. Others are harder to classify. For instance, it is a familiar experience that a full moon seen close to the horizon looks a lot bigger than a full moon seen high in the sky. At first, people thought it was an optical illusion, caused by the fact that light reflected off the moon has to travel through much more atmosphere when it is on the horizon than when it is in the high sky. But if this was so, then one would expect the image on the eyeballs to be of different sizes. But the images are the same size whether the moon is on the horizon or in the high sky. It is now thought that the illusion is produced as the brain “interprets” or processes the information. One possibility is that it has to do with the fact

that the moon on the horizon is seen as close to other objects. But this illusion is not yet fully understood.

4.3 MEMORY

We often rely on our memories to ground or sustain our beliefs. But memory's role in the justification of beliefs is a special one. For memory is not a source of evidence; rather, it is a repository of evidence. Whereas observations are bits of evidence, memories are not. Memory stores evidence. This means that when we rely on our memory, our evidence is no more acceptable than the acceptability of the evidence we remember.

But we also know that memory can be unreliable: it is as if in the storage process the evidence gets modified or changed. The U.S. National Transportation Safety Board (NTSB) provides an especially striking example of this. After a plane crash, the agents of the NTSB collect as much evidence as they can in the hopes of trying to recreate the sequence of events that led to the crash. Among the evidence they collect are reports from eyewitnesses on the ground. But over time, they have found that eyewitness reports are highly variable. In the case of the crash of American Flight 587 in 2001, the NTSB interviewed 394 eyewitnesses. They found that

52 percent said they saw a fire while the plane was in the air. The largest number (22 percent) said the fire was in the fuselage, but a majority cited other locations, including the left engine, the right engine, the left wing, the right wing, or an unspecified engine or wing. Nearly one of five witnesses said they saw the plane make a right turn; an equal number said it was a left turn. Nearly 60 percent said they saw something fall off the plane; of these, 13 percent said it was a wing. (In fact, it was the vertical portion of the tail.)⁴

Why are eyewitnesses so unreliable? It might be that their visual observations of the event are unreliable. But it is hard to see how so many people could have had such different visual experiences, especially when they were all looking at the very same event, and maybe even standing right next to one another. More likely, their visual

⁴ Wald, M. "For air crash detectives, seeing isn't believing." *The New York Times* (2002, 23 June). Section 4, p. 3.

experiences were somehow distorted as they got put into memory, while they stayed in storage, or while they were being retrieved from storage. Whether the distortion happened before the storage, during the storage, or during the retrieval process, their memories are distorted. The NTSB still collects eyewitness reports, but they do this as for largely public relations reasons. They no longer rely on these reports when trying to figure out what happened. Admittedly, memories of horrific visual scenes such as the crash of an airplane are the exception, and the fact that such experiences are misremembered does not by itself show that memory is not in general reliable. But it does illustrate once again the reason for the following maxim: trust, but verify.

BOX 4.3 CRITICAL THINKING IN PSYCHOLOGY

Researchers have found that a person's memory can be manipulated in different ways.⁵ In one study, subjects were shown a fake advertisement of Disneyland featuring Bugs Bunny standing next to the Magic Castle. The ad looked just like a real ad for Disneyland. After studying the ad, subjects were asked to describe their own experience as children visiting Disneyland. Sixteen percent of them said that they remembered meeting Bugs Bunny at Disneyland. The greater the number of exposures to the fake ads, the higher the percent who claimed to remember personally meeting Bugs in Florida. Some even claimed to remember specific details, such as hugging him and touching his ears. But since Bugs Bunny is not a Disney character, these supposed memories are all false, somehow implanted or encouraged by the false advertisements. Researchers have even been able to instill false memories of quite unusual and memorable events. In one study, a subject's parents were enlisted to tell the subject that she had poured a slimy substance onto the head of her grade 1 teacher. The story was the very same for each subject, aside from names and places, and was full of detail. They were even provided with fake photos of the event to show to the subject. Remarkably, 65 percent of the subjects later reported to researchers that they remembered the event in vivid detail, and expressed shock and surprise when told the entire event was fictional.

⁵ Loftus, E.F. "Memories of things unseen." *Current Directions in Psychological Science*, Vol. 13, Number 4, 2004, 145-147.

EXERCISE 4.3

1. List five conditions under which visual perception is not reliable. Do the same for our sense of touch.
2. We learned that our senses are reliable for some topics but not others. List some topics for which vision but not touch is reliable. List some for which hearing and sight are both reliable. List some on which no sense is ever reliable.
3. Sometimes, our different senses provide us with conflicting evidence. Describe such a case. Which sense should we trust in a case like that? If you can, think of a general rule or principle that can be used to always decide which sense to trust when senses conflict.

4.4 TESTIMONY

Suppose that in trying to decide what to believe you only relied on your own observations. How much knowledge would you be able to acquire? Not very much, probably. Just think of how little you can actually see and feel at any one time. Even if our perceptual faculties are as highly reliable as we hope, they are also extremely limited. But we all know (or at least think that we know) a lot about things we have never seen or touched. We know about the history of the U.S. constitution, about the battles of the two world wars. We may have seen the attacks on 9/11 with our own eyes, but we need more than our own eyes to figure out the causes of the attack or what structural causes made the towers fall down. We know about distant places and times, beyond our observation. Most of our knowledge, in fact, would be out of reach if we did not rely for information on other people, whether it be our parents, friends, teachers, authors we read in the news or see on TV shows, or even just people we overhear in the local coffee shop or bar. (Hopefully you have learned a thing or two from me in reading this book.)

Evidence that consists in what other people tell us is **testimonial evidence**. It may sound a bit highfaluting to call the information we get from newspapers, teachers, and parents “testimony.” It sounds a bit odd to say that our best friend is testifying when he tells us that the fridge is full of beer. But the analogy between these ordinary cases of believing what people tell us and the role of witnesses in a trial is very strong. In all of these cases we are treating other people as reliable sources of evidence. We are taking them at their word; trusting their say-so. This raises the question: when is testimonial evidence acceptable?

BOX 4.4A SUMMARY: TESTIMONIAL EVIDENCE

Testimonial evidence is acceptable just in case:

- i. it is on an appropriate topic;
- ii. it is from a competent source;
- iii. the source is unbiased.

The first two conditions are already familiar, though we will see that there are some special factors to keep in mind in the case of competence. But the third condition is a new one. It is needed because whenever we are deciding whether to trust what someone is telling us, we need to think about whether that person is biased or prejudiced. Let's look at each of these conditions in more detail.

4.4.1 Appropriate Testimony

First, testimonial evidence is acceptable only if the topic is **appropriate**. As we saw above, a source of evidence might be reliable on some topics but not others. The same is true for people. When a person is a reliable source of evidence on some topic, we can think of her as like an expert on it. Some people are experts on sports, while others are experts on the chemistry of cells. A person can be an expert on several very different topics. But there are some topics where there simply are no experts. A familiar example is any topic where there are no real objective facts, but just matters of taste. For example, I doubt whether there are facts about whether one popular musician is better than another. There are, of course, people who know a lot about pop music, about the different performers and their histories and musical capacities. There are experts about who can carry a tune and play the guitar. But is there really such a thing as being an expert on whether Madonna's music is better than Bruce Springsteen's? I have always liked Bob Dylan's singing, though many people find it (and my musical preferences) horrifying. But is there anything more to this disagreement than just a difference in taste? I doubt it. I am inclined to think that it is inappropriate to appeal to experts to settle disputes about which pop musician is better.

But it is not just in matters of taste that there are no experts. Sometimes, when a new field of study is very young and just getting established, there will not yet be experts either. This is the case at cutting

edge fields in natural science. When the scientists working at that cutting edge disagree among themselves about the field, and especially when they disagree about which methods are best for measuring or describing the phenomena, then there are probably no real experts yet. In cases where the best-placed people in the field cannot agree, then we as nonexperts should probably withhold judgment too. This was the case at some time in almost every branch of science. It was the case for the science of global warming until about 20 years ago. But now there is no doubt that there are experts on global warming, and that it is perfectly appropriate to rely on what they say when we decide what to believe about global warming.

4.4.2 Competent Testimony

Second, testimonial evidence is acceptable only if the witness is **competent**, and this means only if the witness is both properly trained and properly informed about the case at hand. This is just a way of asking whether the witness is reliable, whether there is a high likelihood that, assuming that he is not biased, his testimony will be true. We can analyze this idea of competence into two elements. First, the witness must have the proper **training and experience**. What it takes to be properly trained depends of course on the topic. Trials provide lots of good examples. Only in special kinds of cases will an eyewitness to a crime have to show that her eyes were working properly the afternoon that she saw the crime, though if she wears glasses that will be relevant to whether she did see what she claims she saw. But it is much more common for a witness on a specialized scientific topic to have to show that she is properly trained. Expert witnesses on DNA testing or finger printing need to show they have the training and certification needed to be acceptable. Usually, the fact that a witness has been certified by the relevant organization is good reason to think she has the proper training to count as an expert witness on that topic. These examples are from court trials. But the same issues arise in more mundane cases too. I would trust my highly trained electrician over my nine-year-old daughter to tell me whether the wiring in the panel is adequate. I would turn to the pharmacist and not the grocery clerk for advice on which antihistamine to buy, though I might trust the grocery clerk over the pharmacist on whether the streetcar will get me downtown fastest. In this last case, it is not the training but the experience that matters.

The difference between being properly trained and being properly informed is an important one. The following analogy might help: a thermometer is really good at telling the temperature of the water in

a glass. But it won't get the reading right unless it is in the water. This is like the difference between being adequately trained and being informed. Sometimes, people who are considered experts find it difficult to admit that they do not know an answer to some question. This is understandable. But it is also an obstacle to critical thinking. We should prefer for them to keep quiet or admit to ignorance than say something that is ill informed. In general, we should not accept what a witness says if we have reason to think she is not sufficiently informed on the issue at hand, even if we think she is an expert on the general topic.

4.4.3 Unbiased Testimony

Finally, testimonial evidence is acceptable only if the witness is **unbiased**. The reason for this condition is pretty straightforward: sometimes witnesses are motivated in different ways to lie about or exaggerate or understate the facts. The example of the murder trial illustrates one possible source of bias. The defense attorneys claimed that the witness was biased because he was being paid by the prosecution to give his testimony. The implication was that the witness might not have given the same testimony had he been paid by the defense attorneys, or by no one. Of course, the fact that the witness was being paid for his testimony does not prove that he was lying or overstating or understating anything. But it might, and for some of the jury it did, raise the possibility that he was biased against the defendant and so was not to be trusted. Desire for financial gain is one source of bias, but it is not the only one. Just about anything can be a source of bias. People are moved to lie by jealousy and by love, by pride and by humility, out of loyalty and out of revenge, by a desire for fame and by a desire for anonymity. If we know that someone is biased, then obviously we should not accept what they say. Their testimony is acceptable only if there is no reason to suspect that they are biased.

BOX 4.4B MISTAKES TO AVOID: UNACCEPTABLE TESTIMONY

It is a mistake to accept testimony from a witness if you have reason to believe that the topic is inappropriate, or that the witness is incompetent or biased. Testimony is appropriate only on topics for which there are recognized experts. An expert must be properly trained and properly informed. And a witness must not be biased—

motivated to lie about, to exaggerate or to understate the facts. This mistake is sometimes called the Fallacy of Appeal to Authority. But it is important to keep in mind that without acceptable testimonial evidence none of us would have very much knowledge at all. Most of what we know (or at least, what we think we know) we learned from others.

Judgments about witness bias can be tricky and can require balancing different facts about the witness. Suppose that the lead scientist for a well-known environmental group testifies before Congress that the water levels in the Great Lakes are dangerously low and that expensive conservation steps must be taken to reduce water usage in the cities and farms that depend on the water from the lakes. What are her motivations? We know that she is paid by the environmental group to champion its environmental policies. If there were no environmental problems to report on, she would be out of a job. She probably also wants to keep her high-profile position, and might enjoy being in the spotlight before Congress. She might think that advocating an extreme position might, given the political realities involved in passing complex regulatory legislation, be the best strategic move. All of these considerations suggest that she might be motivated to exaggerate or even lie about the real findings. On the other hand, it is often hard to get away with a lie, and the reputation of her organization will suffer if it becomes public that its lead scientist has been lying. She also has a professional reputation as a scientist that she probably wants to maintain and even enhance. Lying or exaggerating will surely hurt her image among other scientists. And she probably has a personal sense of honor that forbids lying or exaggerating, except perhaps in extreme cases, which this surely is not. These considerations suggest that she is highly motivated to speak the truth, at least as she sees it. What should we conclude about her testimony? Since we are not experts, we are in a difficult spot.

BOX 4.4C CRITICAL THINKING AND THE LAW

In a famous American trial from the 1990s, the defendant was accused of having viciously murdered his wife and her friend. The prosecution's case was built on forensic evidence that seemed to connect the defendant to blood samples found at the scene of the crime, in the defendant's home, and on a bloody glove of a kind once owned by the defendant. To many outside observers, the case seemed quite

strong. But the defense attorneys did a remarkable job of undermining the prosecution's star witness, a forensic scientist who testified that the blood samples matched the defendant's blood type. First, the defense lawyers argued that this testimony was inappropriate because the science of blood sampling was too young and not yet fully established. Moreover, they argued, even if the testimony were appropriate, the scientist is not competent because he lacked the training needed to use the equipment involved in the sophisticated analysis of the blood. What is more, the defense argued, even if the scientist had been properly training, he was not properly informed because the blood samples that had been collected by the police had been all mixed up and there was no sure way to know which samples were found in which place. Finally, they argued, even if there are experts in this field and even if this witness is both properly trained and properly informed, the fact that he is being paid by the prosecution to give his testimony means that he is biased against the defendant and so is not to be trusted. Point by point, the team of defense attorneys had done a masterful job of undermining the credibility of the witness.

4.4.4 Advertising

Advertisements are a special case of testimonial evidence. Advertisements can serve many purposes, but generating sales of the advertised product or service is surely one of the most important. Advertisements can try to achieve this in different ways. Some appeal to emotions, like the wonderful ads for Apple's iPod music player that involves nothing but a person dancing to the music they're listening to on their iPod. That ad works not by informing us of the product, but by trying to establish an emotional connection between the product and a desirable lifestyle. But many ads do aim to produce sales by informing potential customers about the product or service. Such ads can be thought of as involving testimonial evidence. (Indeed, some even involve "testimonials" by famous people describing their experience with the product or service.) We should evaluate claims made in such ads in the very same ways that we evaluate any other case of testimonial evidence.

Is the claim on a topic for which testimonial evidence is appropriate, that is, on which there are experts? Advertisements for food sometimes include claims about great taste. But this is the kind of topic on which it is at least questionable whether there are experts. Recently, drug companies have been permitted to advertise their products. These ads

are highly regulated and the drug companies are required to provide quite detailed information about potential side effects. But there is an underlying concern about the appropriateness of these claims, given that for many medicines so little is known about potential long-term effects, both positive and negative. This is especially true for claims about the health benefits of diet supplements, since many of those products are neither tested nor regulated by the government, though advertisements of them are subject to the regulations against being misleading that govern all advertisements. In general, the acceptability of claims in advertisements about the health effects or benefits of products is questionable simply on the grounds that such claims are inappropriate.

Advertisements that involve testimonials by famous people raise questions on the grounds of competence. Is a famous movie star really an expert on which phone plan is best for me? (Set aside for now the fact that the actor is being paid to say that it is.) Is the CEO of an automotive company really an expert on the performance of his company's cars relative to the competition? Should we trust what students in college ads say about the benefits of their college compared to the competition? Unless the ad involves a recognized expert there is good reason to not accept the claims made in the advertisement.

Finally, though, given that the advertisement is aimed at producing sales, the risk of bias is inevitable and serious. Advertisers know this and sometimes include favorable evaluations from independent organizations. An ad for a car might refer to the results of crash tests performed by an independent safety group, or might cite awards the car received in independent performance tests. (Of course, claims by outside organizations are themselves just more testimonial evidence and need to be evaluated on their own.)

There are also governmental rules regulating commercial speech, designed to prevent false or misleading ads. But these regulations are difficult to enforce. One recent study, reported in *The New Yorker*, suggests that more than 50 percent of advertisements for nutritional supplements involved false or misleading claims.⁶ The history of advertising also offers little reason to trust claims made in advertisements. It is perhaps best to approach claims made in advertisements with an initial and healthy skepticism: given the high risk of bias, best not to accept the advertised claims.

We have been discussing the conditions under which testimonial evidence is acceptable. We seen that it is acceptable if, but only if,

⁶Specter, M. "Miracle in a bottle." *The New Yorker*, 2 February, 2004, pp. 64-75.

(i) the testimony is on an appropriate topic; (ii) the witness is properly trained and informed; and (iii) the witness is not biased in any way. If any one of these conditions is not met, then the testimony should not be accepted. Of course, testimonial evidence that is not acceptable because it fails to meet these three conditions might still be true. A nonexpert might be right about the facts. A biased person might still be speaking the truth. Deciding that testimonial evidence is not acceptable is not itself reason to believe that it is false. In cases where our only evidence is testimonial evidence and we have determined that the testimony is not acceptable, the reasonable thing to do is to withhold belief.

4.4.5 News Reports

The news media is a further special case of testimonial evidence. News reports, whether in newspapers, magazines, on TV, or on the Internet, all involve a reporter making claims about some topic or other. The reporter might be writing about what happened yesterday on Capitol Hill or in a refugee camp in the Middle East or Africa. Or maybe the report is on economic conditions in Asia or in our local region. Whatever the topic, we should treat what we read or hear on the news in the way we treat other forms of testimonial evidence and we should be ready to ask the same questions. Is the report on an appropriate topic (one on which there are experts)? Is the reporter an expert? Is the reporter biased in any way? Let's look at each question in turn.

Testimonial evidence is **inappropriate** if it is on a topic where there are no experts. This can happen where the facts are so complex that no one counts as an expert. Sometimes, news reports will make claims about the nature of causes of events where it is questionable whether anyone really knows what is going on. It is, of course, not always obvious to us when a topic is that complex. But in some kinds of cases, it is perhaps better to err on the side of caution. Here are two examples. The first concerns reports on the stock market. Reporters who cover the stock market not only report on changes in the value of various stocks as the day goes on, they sometimes offer explanations of why the markets as a whole are moving in one direction or another. The sharp drop was caused by fear that interest rates will go up; or the rise in stocks was a reaction to the morning news that the unemployment rate has once again crept up. These kinds of claims are almost never trustworthy. No doubt there is some explanation of the change in stocks values. But surely that explanation is enormously complex. The same is true, I think, when reporters offer explanations of complex interna-

tional events, and this is especially clear in the case of wars. During the war in Iraq, it was regularly reported that the violence in Baghdad was getting worse in the middle of 2007, even as the Americans sent in more troops. These reports were difficult to assess. It is hard to know how to evaluate violence in terms of worse and better—if there are fewer attacks but each attack is more deadly does that mean the violence is worsening? And because there were so many attacks every day, and because Baghdad is a huge city, it is hard to know everything that is going on. The complexity involved in measuring and defining the level of violence is so high that it might be best to treat claims like that made by reporters as inappropriate.

In general, testimonial evidence is acceptable only if the source is an **expert** on the topic. But this is almost never the case with news reports. Many reporters are trained in journalism schools. This means that they are trained in how to collect information and how to present it in various media. But it does not mean that they are trained in or well informed on the topics on which they report. Indeed, it is surely true that most reporters are not experts on what they are reporting on. This is why they rely on experts in the field when preparing their reports. They present the expert's testimony. In a way, this makes our task as critical thinkers even harder. For not only do we need to assess whether we should accept the reporter's account of the expert testimony, we have also to assess whether that expert testimony is itself acceptable. Cases when reporters rely on expert testimony are like double-testimony! Just to make matters worse, reporters sometimes rely on witnesses who insist on remaining anonymous. The witness might have legitimate reasons to insist on this—perhaps their career or health depends on it. But this makes our task next to impossible: how can we assess whether the witness is properly trained, informed, and unbiased if we do not even know who it is? To some extent, perhaps, we can trust the reporter to tell whether the expert she has interviewed is competent and unbiased. But this is less than ideal.

A news report is acceptable only if it is well **informed**. It is not always easy to tell how hard the reporter worked to collect the information she used in her report. Passive reporting occurs when a reporter merely accepts testimonial or other evidence without doing background checks and without questioning the acceptability of the evidence. We know that most news reports are produced very quickly, and have to be quite short. Articles in the newspaper and stories on the evening news are unlikely to be as well informed as full-length documentaries. The time pressures are too great. But this means that we, as critical thinkers, have to decide whether the report is based on

sufficient information. There are some things to look for in deciding whether a report is well informed:

- i. *Number and variety of sources.* Did the reporter rely on one source or on many? In general, it is better if the reporter asked for many expert opinions. Did the reporter rely on a variety of experts? If all the experts are from a single organization, then there might not be enough variety in their testimony.
- ii. *Background and fact checking.* Did the reporter do any background investigating of her own, or did she solely rely on sources? Without some background work, it can be difficult to know what questions to ask the experts or how to follow-up on their answers. Did the reporter check whether factual claims made by her sources are correct? Or is she simply uncritically reporting what the source told her? The more reason there is to think the reporter relied on a large number and variety of sources and did some background and fact checking, the more acceptable the report will be.

Like any testimonial evidence, a news report is acceptable only if it is **unbiased**. In the case of news reports, bias is a potential factor at several different levels, starting with **reporter bias**. Reporters are under pressure to produce, just like anyone else. And we all know that salacious and juicy stories are more fun to read than dry and factual stories, even if the factual ones are intrinsically more important or newsworthy. “If it bleeds, it leads.” Reporters are under pressure not just to report the news, but also to report what they think their audience wants to read or hear about. Very rarely, reporters even react to this pressure by making up the news, or by focusing their report on what they think will be most interesting or catchy. **Corporate bias** is also a potential factor. Newspapers and news stations are businesses and this means that they are in the business of making money. They do this by reporting on what their audience is interested in. Inevitably, this means that they leave some stories completely unreported. The latest celebrity arrest for drunken driving is covered in more detail than the thousands of children who die every day in refugee camps around the world. The fact that news organizations are in the business of making money also means that they are reluctant to report stories that would make their audience feel uncomfortable. A newspaper that published every day the names of all the children who die of hunger would not be very successful. News organizations are also careful not to offend the patriotism of their audience. News reports that focus on war crimes sometimes

face criticism. Finally, news reports are subject to **cultural biases**. This affects not just what stories are told but how they are told. We find stories about local events more interesting than stories about events in distant parts of the world. It is not easy to detect bias. Respectable news organizations work hard to draw a sharp line between news reporting and editorializing, and some have ombudsmen whose job it is to keep an eye on and even report on the extent to which the organization is succeeding at being unbiased. Finally, it is always good advice not to rely on only one source for news. Reading different reports on the same event is the best way to avoid falling prey to biased reporting.

BOX 4.4D MISTAKES TO AVOID: UNJUSTIFIED CRITICISM OF TESTIMONY

We are justified in accepting testimonial evidence unless we have good reason to think that it violates one of the conditions. (The acceptability of claims in advertisements is a special case since the presence of bias is inherent to advertisement.) It is the **absence** of reason to think that the claims are unacceptable that matters. I can accept testimonial evidence even if I do not have evidence that the witness is competent or unbiased. But the case is different when we decide not to accept testimonial evidence. The mere possibility that the witness is incompetent or biased or that the testimony is inappropriate is not good enough reason not to accept the testimony. It is not enough for the defense attorneys to charge the expert witness with being biased. The defense attorneys have an obligation to show that the witness' testimony is in fact biased. Likewise, you need to have good reason to not accept what someone says.

EXERCISE 4.4

1. Short-answer questions:

- a.** Under what conditions is testimonial evidence acceptable?
- b.** What is the difference between being trained and being informed? Illustrate your answer with an example.
- c.** If a witness is biased, does this make their testimony false? Give an example to illustrate.
- d.** If a witness is testifying about his own personal observations, what are the critical thinking questions that we should ask before accepting his evidence?

- e. Under what conditions is testimonial evidence overridden and undermined? Use examples to illustrate your answer.
 - f. What are some sources of media bias?
 - g. What is passive reporting? Is it a form of bias? Why or why not?
 - h. When are claims made in advertisements acceptable?
2. The following passages involve appeals to testimonial evidence. Determine whether the evidence is acceptable. If not, then identify which of the conditions is violated. Be as detailed as you can.
- a. One of Thomas Jefferson's most trusted advisors said that the U.S. should not trade with tyrants. We should take that advice and cut off all economic relations with tyrannical regimes around the world.
 - b. There is milk in the fridge. I just called home and Joan told me that there is.
 - c. I am failing this class. My teacher just told me so.
 - d. The cookbook says to boil the eggs for 12 minutes to make them hard, so this is what I am doing.
 - e. The newspaper just reported that the stock market will drop tomorrow, so I'm selling all of my stocks now.
 - f. The bank president says that there is no risk that his bank will go out of business, so I've decided to keep my money in it for now.
 - g. The man at the garden supply store told me that this plant will thrive best in a shady spot, so I'm going to put it underneath that tree.
 - h. The man at the garden supply store told me that this hose won't leach lead poisoning into the soil. That's why I bought it.
 - i. I think that you will really like this band. The girl at the music store said that they are the best all-boy band since 'N Sync.
3. The following passages involve attempts to undermine a witness' testimony. Using the concepts we have discussed in this chapter, explain whether the attempt is successful. If not, explain why it fails. Be as detailed as you can.
- a. The man at the garden supply store said that the plant thrives in shade, but he just stocks the shelves there. I do not trust him.
 - b. Our local politician says that we should revise campaign finance legislation to make it harder for politicians to be influenced. But he's just trying to raise funds from the "little guy." Don't trust him.

- c. The regional coordinator of that environmental group gets paid to recruit new members, so we cannot trust what he says about the effects of global warming. He's just trying to get us to contribute.
 - d. The newspaper article says that the new highway is very dangerous. But I do not think the reporter has any background in that field. He cites a lot of expert reports, but he has no training himself. We should not trust him.
 - e. The witness claims that Jones robbed the bank, and that she saw him leaving with the bag of money. But there were so many people coming and going that day that there is no way the witness could be reliable on this. No one could remember every face they ever saw.
 - f. The cookbook says that it is best to add fresh anchovies to the dressing, rather than canned ones. But the cookbook author owns a chain of fresh seafood stores, so she's probably just trying to increase sales.
4. Look at a newspaper article. Using the concepts discussed in this chapter, assess whether the reporter is guilty of passive reporting. Be as specific as you can in your criticisms, and make sure to support them carefully.
 5. Look in the newspaper or on TV for five examples of testimonials. Assess whether the testimonial evidence is acceptable. Be as specific as you can, and be sure to support your conclusions with reasons.
 6. Look at today's newspaper and find a report on some international incident. Then go to the library or go online and find news articles from different continents on the very same incident. Compare the reports looking for significant differences.

4.5 MEASUREMENT

When we try to decide what to believe or what to do we often rely on evidence from measurement. After all, one way to collect information on some subject or phenomena is to measure it. We have a huge number and variety of measuring instruments and tools at our disposal. We measure our own mass using a bathroom scale and that of objects in deep space using highly sophisticated instruments; we measure public opinion using surveys and questionnaires; we measure the intelligence of our children using IQ tests; meteorologists measure the speed and

direction of traveling storms; we measure student performance using final exams. (What, in your view, do final grades actually measure?) Businesses use personality tests to find out about their employees and to build better teams. It would not be much of an exaggeration to say that measurement plays as important and prevalent role as observation or testimony in our reasoning about what to believe and do. Thinking critically about the acceptability of measurements involves asking some of the very same questions we asked about observation and testimony.

Here's an initial question: what is it, exactly, to **measure** something? Measuring usually involves assigning a number. But maybe not always: Doctors sometimes ask children with sore throats to indicate the level of soreness using a chart of face drawings, ranging from a happy smiling one to one that is crying. Are the kids using this chart to measure their pain, even though no numbers are involved? And assigning a number to something is not always measuring it. Counting, for instance, involves assigning a number. We count and find that there are five donuts on the table. Have we measured the donuts? We have not measured their size or mass. Have we measured their number? This sounds odd. We do, though, measure the size of crowds.

Calculating averages also illustrates the difference between measuring a phenomena and merely assigning it a number. As you may know, there are many different kinds of averages. Suppose that 100 students wrote a final exam and that each was assigned a grade from 0 to 100. The **mean** grade is the result of adding the 100 grades and dividing by 100. The **median** grade is that grade such that half of the grades are above it. The **mode** grade is the most common grade. These three averages vary independently of one another—changing any one of them might not change the other two. Which kind of grade measures student performance depends on what question we are asking. More students might be in the B-range, even if the mean is C– and the median is a C+. The moral is clear: simply assigning a number to a phenomenon is not the same as measuring it.

When we measure something we assign a number and relate that number to a standardized scale of some kind. We say, not just that my weight is now 175, but that it is 175 pounds. The concert lasted not just 30, but 30 minutes. The restaurant review gives the new Indian restaurant not just 5, but 5 stars. Understanding a measurement requires knowing which scale is being used. The acceptability of a measurement depends on the scale too. Some scales are appropriate for measuring some aspects of a phenomenon but not for measuring others. We can measure a liquid's temperature but not its mass using degrees Celsius.

We can measure a car's velocity, but not its acceleration, using meters per second. In some cases, there is more than one appropriate scale, as in the case of temperature, which we can measure using either the Celsius, the Kelvin, or the Fahrenheit scale. So, one question we need to ask in deciding whether some measurement is acceptable is whether the measurement employs an **appropriate scale**.

We measure something using a measuring instrument of some kind. We measure our mass using a bathroom scale; temperature using a thermometer; student performance using a final exam; voltage using a voltmeter; public opinion using a questionnaire. Whether a measurement is acceptable depends on whether the instrument used in the measurement tends to yield accurate measurements. A measuring instrument is **reliable** only if it tends to accurately measure what it is supposed to measure. If a test is designed to measure a person's latent hostility, then it is a valid instrument only if it does in fact provide accurate measurement of the latent hostility of people who take the test. A procedure for measuring the size of crowds at rallies and demonstrations is valid only if it provides an accurate (or accurate enough for the purposes at hand) measure of the crowd size.

It is sometimes difficult to know for sure whether an instrument is reliable. Just because something is called an intelligence test does not guarantee that it really is a reliable instrument for measuring intelligence. Likewise, just because something is called a public opinion survey, this does not mean that it reliably measures public opinion. (We will study the use of such surveys in more detail in Chapter 6.) There is a deep and difficult methodological problem involved in determining whether a measuring instrument is valid. The only way to know for sure whether a measuring instrument is valid is to compare its readings with those of a device that is known to be valid. This is called "calibrating" the instrument. So, I can check to make sure that my meat thermometer is still valid by comparing its readings to one I know is valid. But what if we do not know whether that second device is valid? Do we need to compare its reading with those of a third instrument? Where would this regress end? Or what if we do not have another instrument for measuring the phenomena? If we have independent access to the phenomena being measured, then this is not a serious problem. Measuring the length of wooden boards is such a case, since we can more or less confirm by looking whether the tape measure is giving us the right reading. But with tests used to measure human intelligence or student performance, this is a serious problem, since we often have no independent way to measure the phenomena. We have no independent way to check to make sure that

it really is the students' intelligence, as opposed to some other trait, that is being measured.

Thankfully, we do not need to sort out this difficult problem here. It is enough for us to be aware of the methodological problems facing measurement, and to be armed with the concepts needed to think critically about them. But in deciding whether a measuring instrument is valid, there are two very important points to keep in mind here.

4.5.1 Measurement Consistency

First, consistency is not a guarantee of reliability. An instrument is consistent if it gives the same readings on repeated uses. Whether a reliable measuring instrument must be consistent depends on the nature of the phenomena being measured. If it is a phenomenon that can change quickly between measurements, such as levels of sugar in a person's blood or the electricity usage of a building, then a reliable instrument might not be consistent, since its readings would have to change to keep track of the changes in the phenomena. An instrument that always gave the same reading of a person's blood sugar levels throughout the day would be very consistent, but that might actually be an instrument that it is not reliable, since a reliable instrument for measuring blood sugar should give different readings at different times of the day. On the other hand, if the phenomenon does not change rapidly between measurements, then a reliable instrument would have to be consistent. If my bathroom scale gives five different readings within a few minutes then this is good reason to think it is no longer reliable. If a drugstore pregnancy test gave different readings every five minutes, then we would have good reason to question its reliability. So, whether consistency is a virtue in measurement depends on what it is that we are measuring. An inconsistent instrument might be reliable, and a consistent instrument might be unreliable.

If we know that the phenomenon we wish to measure does not change very quickly, then a measuring instrument would have to be consistent in order to be reliable. The easiest way to tell whether an instrument is consistent would be to use it to measure the same case several times. I can test the consistency of my new food scale by repeatedly putting the same bag of potatoes on it to see whether it always gives the same reading, which it would if it was consistent. But just to make matters a bit more complicated, in certain kinds of cases, an instrument designed to measure a very stable phenomenon might nonetheless not be consistent. This is the case for instruments designed to measure human intelligence or knowledge, like the SAT test or

standardized tests in grade school. A subject who has taken a standardized test once knows all the questions and so it is likely that she would do better the second time around. It is not that taking the test has by itself made her smarter or more knowledgeable (although she does know more about the SAT test) if we did repeat the test, we would expect different outcomes. But this need not be reason to think that the test is not reliable. The relations between instrument reliability and instrument consistency are obviously quite complex.

4.5.2 Measurement Precision

Second, precision is not a guarantee of reliability. The precision of an instrument is a matter of how finely graded its readings are. A scale that measures only in pounds is less precise than one that measures in ounces. But precision is no guarantee of reliability. If the more precise scale is off by more than two pounds, then the less precise one is more reliable. Still, precision is alluring. Sometimes, the results of public opinion polls involve decimal points, such as that 56.3 percent of the population is opposed to some policy. This precision seems to suggest that the instrument they used to measure public opinion must be very accurate. But this is a mistake. Precision is neither necessary nor sufficient for reliability.

BOX 4.5.2A DECIDING WHAT TO DO: MEASURING COSTS AND BENEFITS

Deciding what to do requires comparing the anticipated costs and benefits of competing proposals. But this comparison is often very difficult. Here are two of the main reasons for this:

Incomplete information. Sometimes, it is just not possible to know what the anticipated costs and benefits of a given proposal are. Unfortunately, global warming is a good example. The nature of the Earth's climate is so complex and its dynamics so subtle that it is next to impossible to know which of the many proposals for counteracting its effects will be most efficacious.

Incommensurability. Sometimes, the costs and benefits of a proposal resist the kind of quantification that is essential to measurement. Decisions that affect quality of life are like that. There is no easy way to compare and contrast pains and pleasures.

Inevitably, we sometimes must decide on a course of action in the absence of complete information or in the presence of incommensurability. In such cases, it is good to make the ignorance or incommensurability as clear as possible. If one must act in the face of ignorance and incommensurability, it is better to do so knowingly than blindly.

BOX 4.5.2B MISTAKES TO AVOID: CONFUSING THE LIKELIHOOD AND THE VALUE OF A COST OR BENEFIT

It is a mistake to confuse how good or bad some consequence would be with how likely it is. Winning the lottery would be terrifically good, but it is extremely unlikely. So when comparing the costs and benefits of alternative proposals, you need to factor in the likelihood of those costs and benefits, as well as their value. A course of action that promises very high value, but at very low probability (e.g., spending your money on the lottery), might not be as good as one that promises a high probability of moderate value (investing cautiously for the long run). This distinction between the likelihood and the value of a cost or benefit is especially important to keep in mind when assessing risk. Being killed in a plane crash would be far worse than getting a flat tire on the highway, but the odds of the flat are much higher than those of the crash. Flat-tire insurance is a better investment than crash insurance.

BOX 4.5.2C MISTAKES TO AVOID: APPEALING TO IGNORANCE

It is a mistake to discount or—what is even worse—deliberately ignore potential costs or benefits of a proposal on the grounds that the costs or benefits are either unknowable or incommensurable. Doing so would in effect be appealing to ignorance.

4.5.3 Surveys

Opinion surveys are a very familiar form of measurement. Pollsters and researchers ask people for their opinions on everything from politics, to sexuality, to sports, to food, to history.

BOX 4.5.3A SUMMARY: WHEN IS EVIDENCE FROM A SURVEY ACCEPTABLE?

Evidence from a survey is acceptable just in case:

- The survey questions are not ambiguous, biased, loaded, or otherwise bad;
- Those surveyed are properly trained and informed on the topic of the survey;
- There is no evidence of researcher or subject bias.

Researchers even use questionnaires to learn more about the nature of happiness, as we will see in a moment. Often, the researchers are interested in measuring the opinion of a large population—perhaps all Americans—and administer their survey to a sample of the general population. In Chapter 6, we will consider when reasoning using samples is valid, but here we will stay focused on the question of when the results of a survey are themselves acceptable.

It pretty much goes without saying that happiness is a prime motivator in our life. Everybody wants to be happy and most of us want to help make others happy too. Companies and governments want people to be happy with the products and services they provide. Economic theory assumes that we are all “happiness maximizers,” that we can be expected to act in ways that we think will make us the happiest. If we choose the burger over the green salad, then that shows that we believe that we would derive more happiness from the burger. In moral theory, Utilitarians hold that the moral value of an action, practice, or policy is a matter of how much overall happiness that action, practice, or policy would produce as compared to alternatives. Public discussions about whether to implement one kind of government program or another often turn on questions about which program would have the best results, where this is usually a matter of how much happiness it would produce. Building a new bridge will make truckers and commuters happier, but it will make those living near the waterfront less happy. Deciding whether to build the bridge requires thinking hard about happiness. Given the central role that happiness plays in our thinking about our own lives and in our thinking about public policy, it might come as a bit of a surprise that we actually have little idea how to measure it.

Suppose that we were trying to measure happiness by asking people a simple question like this: “On a range of one to ten, how happy would you say that you are?” The question is simple enough that most people would be able to answer it. It includes a numerical range so we can get a number to deal with. If we asked enough people, we could even calculate an average of some kind (a mean or a mode or a median?), and figure out how happy the average person is. We could do the same for people under 30 and college kids and grandparents, and thereby learn quite a lot about how happiness varies from one group to another. But in fact there are good reasons for doubting whether this simple instrument is really a valid instrument for measuring human happiness.

One difficulty is that the question is a **bad question**. There are (at least) two problems with the way it is worded. First, it is **ambiguous**. Opinions about what it is to be happy vary quite widely. It is not just that people find that different things in their life make them happy. Some people find that a career-centered life makes them happy, while others derive happiness from their hobbies or from a vibrant community of family and friends. Variation from one person to another—and from one culture or time period to another—in views on the essential elements of a “good life” is not at all surprising. But people who agree on the elements of a happy life might still disagree about what it means to be happy. Is happiness, as John Stuart Mill thought, a kind of feeling, like an emotion or a sensation? If being happy is the same as feeling pleasure, then maximizing happiness is a matter of maximizing the amount of pleasure in the world. Or is happiness, as Aristotle thought, a matter of achieving and sustaining a kind of balance in one’s life, independently of one’s actual feelings? In that case, maximizing happiness would require organizing the elements of one’s life—work, family, friendship, love, sport, pleasure, etc.—in a special way. Of course, people who agree on this view of the nature of happiness might also differ over just what elements are required and how to properly balance them. Disagreements over the nature of happiness are, as we know from Chapter 2, disagreements over the meaning of the word “happiness.” So, the question is bad because it is ambiguous.

The other reason to think the question is badly framed concerns the point we saw earlier that measuring something is not the same as simply assigning it a number. Measuring requires both a number and a scale. But what is the scale in our proposed question? What are the units? Happiness units? But what are they? Moreover, what is the relation between the numbers on the scale? Is it like the relation between numbers on a scale for measuring mass, where something that weighs six pounds is twice as heavy as something that weighs only three

pounds? Is a person who is 3 on the happiness scale half as happy as one who is a 6? Or is the relation between the numbers on the happiness scale like the relations between numbers on a temperature scale, where 20 degrees Celsius is not twice as hot as 10 degrees Celsius? Because we do not know the answers to these kinds of questions, we do not know what the question is asking us. This is a second reason to think that our proposed survey question is a bad question.

There are other ways that a survey question can be bad. A question is bad if it uses **charged** or **slanted** words. If we want to know a person's opinions about the morality of abortion, it won't do to ask whether they are in favor of the deliberate killing of unborn humans. A question is bad if it presupposes something such that no matter how one answers it one will be saying or implying something controversial. The classic case of "Have you stopped cheating on your exams?" nicely illustrates the problems with a **loaded** question. Answering "Yes" implies that you once did cheat; answering "No" implies that you still are cheating. Lawyers questioning a witness in a trial have to be very careful that their questions do not presuppose anything controversial. The goal is always to get as good a handle as possible on the truth, and badly framed questions are an obstacle to the pursuit of truth.

But even if we could agree up front on the nature of happiness and on the scale for measuring it, there are other reasons for doubting the validity of that simple survey question as a measure of human happiness. Because the survey asks people to report on their own level of happiness, it is a measuring instrument that relies on testimonial evidence. As we have seen, testimonial evidence is not acceptable if we have reason to think the witness is **incompetent** or **biased**. So the instrument is reliable only if people are pretty good at telling whether they are happy and only if they are not motivated to exaggerate or lie about their levels of happiness. As it turns out, we have some reason to think that people are not very good at detecting their own levels of happiness. We are all familiar with the way that people live in denial. People sometimes lie to themselves about their true feelings, fabricating stories about themselves and their lives to paper over the problems that lie just beneath the surface. And if happiness is a matter of having a well-balanced life, it is likely that some people will not be good at telling when their life is well balanced. Researchers have found that people's reports on their happiness are not very consistent and can be influenced by irrelevant factors. In one study, researchers found that subjects tended to report a higher degree of happiness if the sun was shining or if they had just found a bit of money (left on purpose by the researcher). Subjects also focus on the best events and ignore or fail to remember

the negative events in their daily life. All of this suggests that some amount of error and ignorance about one's own happiness is inevitable. If so, then we have some reason to doubt whether people are good at telling whether they are happy, and this is some reason to doubt the validity of the proposed measuring instrument.

Surveys of popular opinion on policy matters also raise questions about competence. It is one thing to have an opinion, but quite another to have an educated or justified opinion. Consider a poll asking for opinions on a certain plan for changing the health care system. Most people pay very little attention to the details of government proposals. And most of those proposals are so complex, and most of the problems they are nominally intended to remedy are so multi-dimensional that it is not very likely that most people will have an educated opinion on the matter. Even if the survey accurately reports people's opinions, these opinions are themselves of only little value since they are probably not educated opinions.

BOX 4.5.3B CRITICAL THINKING AND HUMAN HAPPINESS

Researchers on human happiness have proposed a different series of questions to measure a person's happiness. Instead of asking something like our proposed question, some have used what is called a "day reconstruction method," where the subject is asked to describe the sequence of events in their day and report on the emotions they felt during each event. These researchers hope that this method will lead to a more accurate report from the person about just how happy they are with their day-to-day life. So there may be ways to formulate surveys that will enable more accurate reports from subjects. (For a discussion, see "The not so dismal science: how economists measure happiness," Tim Harford, *Slate*, <http://www.slate.com>.)

But suppose that people were in fact quite good at telling whether they are happy. Their testimonial reports would still be unacceptable if we had good reason to think that they are **biased**. The possibility that they are biased would be a case of what in critical thinking we call **subject bias**. It is probably not right to suspect that most people would lie about how happy they are. We do know that people tend to exaggerate what economic class they are in, so that people well in the working class tend to report that they are in fact in the middle class. So it is probably right that we should expect that some people will be motivated to overstate

their happiness, either out of shame or in order to impress. There is also in the case of this kind of survey, the possibility of **researcher bias**. The researchers asking the questions might themselves introduce a bias into the reports, either in the way they ask the questions or in other subtle features of their interactions with the subjects.

The best way to eliminate the risk of subject or researcher bias is to make the test **double blind**. A study is double blind if both the subjects that are being studied and the researchers doing the studying are ignorant of key facts about the study. In the case of drug trials, researchers ensure that their study is double blind by ensuring that neither the research subject nor the researcher knows who got the placebo and who got the trial drug. One way to reduce the risk of bias in a survey is to add a lot of questions on irrelevant topics. That way, neither the subject answering the questions nor the person hired to ask the questions knows what the researchers are really trying to find out about. This ignorance reduces the risk that bias will distort the results.

EXERCISE 4.5

1. Short-answer questions:

- a. What is the difference between mode, median, and mean? Use examples to illustrate when each would be valuable.
- b. What is the difference between measurement validity, measurement consistency, and measurement precision? Use examples to illustrate your answer.
- c. Suppose that you had a watch that was designed only to tell the hour and another designed to tell the time to the closest second. Which is most consistent? Which is most valid? Which is more precise?

2. Make a list of ten measuring instruments you use every day. Order them from the most valid to the least valid. Order them from the most precise to the least precise. Order them from the most consistent to the least consistent.

3. Make a list of five phenomena that you do not currently know how to measure, but which you think should be measurable. Pick one of them and think of a way to measure it.

4. Look through the newspaper and find five articles that make measurement claims. Using the concepts we discussed in this section, identify some questions that need to be considered in deciding the measurements are acceptable.

5. In the following passages, identify the premises and conclusion(s). Then identify whether the argument relies on evidence from observation, testimony, or measurement. Finally, using the concepts we have discussed in this chapter, identify some questions that need to be considered in deciding whether the argument's premises are acceptable. Be as specific as you can.
- a. The roast is probably done. I just took out the thermometer and it read 150 degrees, and the cookbooks says that a roast is done when it reaches 145 degrees. It also smells like it is starting to burn.
 - b. It is safe to go into the water. The city tested the water yesterday and the level of potentially infectious chemicals was very low.
 - c. The new standardized math test is really accurate. We used it on several students from grade three and they all got the same score.
 - d. You have heard the witness testify that she saw the defendant enter the bank on the time of the crime. And you have been presented with ballistics evidence proving that the defendant's gun was used in the shooting. You have no choice but to find the defendant guilty.
 - e. Our new American Motors sedan is the safest car in America. It scored a record 98.79 overall safety rating in our crash tests, the highest of any car we have ever manufactured. Just watch what happens in this video of a head-on car crash. See how the airbags inflate in time to prevent serious injury. This shows that the car is as safe as can be.
 - f. Beautyderm skin lotion removes 87.95 percent of wrinkles. Our clinical tests show it.
 - g. South Park High School is the best school in the city. Its students recently scored higher than every other student in the city on the new standardized tests.
6. For each of the arguments in the previous question, decide whether the premises provide sufficient support for the conclusions. If not, suggest ways the argument might be strengthened.
7. Look in the newspaper for articles reporting on surveys. Find the questions that were asked, and assess whether they are ambiguous, loaded, biased, charged, or otherwise bad. Be as specific as you can, and do not charge a survey with asking a bad question unless you can support your claim.

BOX 4.5.3C SUMMARY

Reasons to believe or do something are acceptable when they are either sufficiently supported by acceptable reasons or when they are from a **reliable** source and there is no **undermining** or **overriding** evidence. A source of evidence is reliable if it tends to provide true or accurate information most of the time. Observation, testimony, and measurement can be reliable sources of reasons.

4.6 MISTAKES TO AVOID

Confusing Undermining and Overriding Evidence It is a mistake to think that evidence from an unreliable source must be false. When we have evidence that some source is unreliable, then we have good reason not to accept that evidence. But it might still be accurate. Just because the witness is not properly trained or is biased against the defendant, this does not mean that her testimony is mistaken.

Testimonial Fallacy it is a mistake to accept testimony from a witness if you have reason to believe that the topic is inappropriate, or that the witness is incompetent or biased. Testimony is appropriate only on topics for which there are recognized experts. An expert must be properly trained and properly informed. And a witness must not be biased—motivated to lie about, to exaggerate, or to understate the facts. This mistake is sometimes called the Fallacy of Appeal to Authority. But it is important to keep in mind that without acceptable testimonial evidence none of us would have very much knowledge at all. Most of what we know (or at least, what we think we know) we learned from others.

Measurement Fallacies It is a mistake to confuse the validity of a measuring instrument with its consistency or its precision. An instrument might be valid even though it is neither consistent nor precise; and neither consistency nor precision is sufficient for validity.

Confusing the Value and the Likelihood of a Cost or Benefit It is a mistake to confuse how good or bad some consequence would be with how likely it is. Winning the lottery would be terrifically good, but it is extremely unlikely. When comparing the costs and benefits of alternative proposals, you need to factor in the likelihood of those costs and benefits, as well as their value. A course of action that promises very

high value but at very low probability (e.g., spending your money on the lottery), might not be as good as one that promises a high probability of moderate value (investing cautiously for the long run).

Fallacy of Bad Question It is a mistake to ask a question that is ambiguous, contains charged or slanted words, or that hides a controversial presupposition.

4.7 PRACTICAL STRATEGIES

Trust, But (Be Prepared to) Verify Most critical thinking theorists agree that it would be asking too much to require that before the evidence from some source can be accepted we must first determine whether the source is reliable. Instead, they recommend the following: evidence from some source is acceptable; unless one has reason to think the source is not reliable. Trusting our sources is a default right, as it were. But we should not let ourselves get carried away. For we know that some apparent sources are not reliable at all, and others even ones that are reliable can still yield mistaken evidence. To borrow Ronald Reagan's remark about the proper attitude to take to enemy super-powers: trust, but be prepared to verify.

Measure Twice, Decide Once The goal of critical thinking is knowledge, and this means that it places a premium on getting the right answer. To this end, it is better (to paraphrase a familiar wood-working lesson) to measure twice and decide once. This applies just as much to observation and testimony as it does to measurements. It is, all things considered, more prudent to collect more evidence than less evidence, and evidence from different sources is best of all. Finally, it is just as important to consider possible undermining evidence as it is to consider possible over-riding evidence. For part of what makes critical thinking reflective is that it requires us to think about what makes a source of evidence reliable.

4.8 FROM THEORY TO PRACTICE: APPLYING WHAT WE HAVE LEARNED

Thinking Critically about Ourselves This exercise is designed to help you reflect on your strengths as a witness. In Chapter 1, you compiled a list of five or six character traits that you think are essential to being

a morally good person. Being a trustworthy person might not have been on that list, but I think that most of us would agree that we strive to be someone others can trust. We have seen that testimonial evidence is acceptable only when (i) it is appropriate; (ii) the witness is properly trained and informed; and (iii) the witness is not biased in any way. This exercise is designed to have you reflect on the extent to which you meet these criteria. As always, the more sincere effort you put into it, the more you will get out.

For two days, observe yourself as you answer people's questions or give them information, or tell them your beliefs. As you do this, be willing to think critically about whether you are meeting the standards for being an acceptable witness. Do you ever offer an opinion as if it were the truth on a subject where there may not be experts? Do you ever offer a firm opinion on a subject where you are not really fully trained or fully informed? Do you ever let biases creep into your responses?

One way to do this exercise is to keep a journal for two days, pausing at noon and before bed, to reflect and describe a few events from the day. Use the Testimonial Rubric we discussed in the text.

Thinking Critically in the Classroom This exercise is designed to help you identify the different sources of evidence that you rely on in studying or engaging in your chosen discipline. Some disciplines rely on one source of evidence much more than others. (Philosophy involves virtually no measurement at all, and little direct observation.) In Chapter 1, you compiled a list of five or six of the tasks that you are required to do in your chosen discipline that require critical thinking. That list might have included such things as performing measurements, collecting observational data, and doing factual research. Make a list of the kinds of evidence that you rely on in studying or that someone actively engaged in your chosen discipline would regularly rely on:

- i. Direct observation
- ii. Reliance on testimony
- iii. Measurement

Thinking Critically at Work This exercise is designed to help you think critically about the sources of information you or the organization you work for rely on for success. All organizations rely on testimonial evidence and measurement in order both to achieve their organizational goals but also to make adjustments to their internal operations. In Chapter 1, you compiled a list of tasks that you regularly

perform at work that require critical thinking. Look over that list, and identify those tasks that require you to collect or rely on testimonial evidence or measurement. Pick one of each kind, and do the following:

1. Assess how reliable those sources typically are.
2. Think of some ways to improve their reliability.
3. What other testimonial sources or measuring instruments would help you with your task?

5

REASONING ABOUT ALTERNATIVES AND NECESSARY AND SUFFICIENT CONDITIONS

Critical thinking is reasonable and reflective thinking aimed at deciding what to believe or what to do. In Chapter 1, we saw that thinking critically requires having or finding good reasons for the decisions we make and the beliefs we form, and we saw that reasons are good ones only if they are both acceptable and sufficient. Since critical thinking is aimed at knowledge, the kinds of reasons it requires are epistemic reasons, which are more commonly called “evidence”. In Chapter 2, we learned how to think critically about the meaning of claims, plans, and proposals. In Chapter 3, we studied what it is for our evidence to be **sufficient** and we saw that this requires that it provide enough of the right kind of evidence. In Chapter 4, we saw that our evidence is **acceptable** when it comes from a reliable source that we have no good reason to doubt, or when it is itself sufficiently supported by acceptable evidence. If in our thinking, we rely on sufficient and acceptable evidence for beliefs and plans that are as clear as we can make them, then we have pretty much done our duty as critical thinkers. Though our beliefs might still be mistaken and though our plans may yet fail, we will have done the most we can in advance to avoid this. We will have pretty much fulfilled our intellectual obligations as critical thinkers.

But it would be nice to have easier ways to tell when our reasons are sufficient. In Chapter 3, we studied the notion of validity, which is the ideal level of support. An argument is **valid** just in case it is not possible for the evidence contained in the premises to be true and yet for the conclusion to be false. That means that the evidence provides the very best kind of support—it leaves no room for error, at least in the sense that if the evidence is accurate or true, then the conclusion is guaranteed to be true too. We learned a method for testing whether an argument is valid: suppose for the sake of the argument that its premises are true, and then ask whether there is any conceivable way that its conclusion might nonetheless be false. If the answer is NO, then the argument is valid. But we also saw that this method is tricky to use. It requires a healthy imagination to suppose for the sake of an argument that the premises are true and then to consider whether there is any conceivable way for the conclusion yet to be false. Thankfully, there are some forms of argument that are guaranteed to be valid and for this reason are very commonly used in deciding what to believe or what to do.

BOX 5.1 PRACTICAL TIP: TESTING FOR LOGICAL STRENGTH

To test the logical strength of an argument, suppose that the premises were true. Then ask how likely is that the conclusion would be true too? The higher the likelihood, the more logical support the premises provide.

In this chapter and the next, we will study three very familiar and very useful forms of reasoning, methods for drawing conclusions from the evidence that we can rely on when thinking critically about what to believe and what to do. In this chapter we will look first at reasoning about alternatives and then study reasoning about necessary and sufficient conditions. In the next chapter, we will examine reasoning using analogies. As you will see, you have probably already relied on these forms of reasoning in your own thinking. Indeed, you are probably already pretty good at that kind of reasoning. Still, by studying their different natures and by identifying some of the more common mistakes that can be made, you can help make your reasoning more reflective. Being reflective in one's reasoning is a core element in being a critical thinker.

5.1 REASONING ABOUT ALTERNATIVES

Often when we try to decide what to believe or what to do we are faced with several alternative possibilities from among which we have to choose and we reason to a conclusion by trying to rule some of them out. I don't remember where my watch is, but I do know that it is either in the bathroom or in my backpack. The city can either build a new bridge or renovate the one it has, and has to find a way to decide between these options. The patient's symptoms are consistent with both a simple viral infection and a more severe early form of cancer, and the medical team needs to figure out which it is. Jones can continue his education by going to either law school or medical school, or he can follow his bliss by pursuing an acting career. Sherlock knows that either the maid did it or the butler did it or the chauffeur did it, but he is not sure which. Before he reaches a conclusion about who the murderer is, he needs to rule out some of the possibilities. Reasoning about possibilities is one of the most common forms of reasoning, and it is important to do it well. In this section, we will study the nature of this kind of reasoning and identify some mistakes to avoid.

5.1.1 The Meaning of Disjunctions

When we reason about alternatives, we usually formulate them using a **disjunction**. As we saw in Chapter 3, a disjunction is just a statement containing the word "or." (It might be worth taking a moment to review that section before reading on.) Here are some examples of disjunctions:

Jones can go to law school or to medical school.

The patient either has a viral infection or an early form of cancer.

Either the sickly plant needs more water or else it has been overwatered and needs to be allowed to dry out.

The defendant either lied before the grand jury or else she told the truth.

Using the Assertion Test, we can see that in asserting a disjunction one does not assert either disjunct. In claiming that the patient has either a viral infection or a form of cancer, the doctor is not claiming that the patient does have cancer and nor is she claiming that the patient has a viral infection. All she is claiming is that an infection and cancer are the only two possible causes of the patient's symptoms. Perhaps she has

already decided which of the two she thinks is most likely. But in asserting that disjunction she has not yet said that. As we saw in Chapter 2, a disjunction is true just in case at least one of the disjuncts is true. This means that the doctor's assertion is true just in case one of the two possibilities she described is the real cause. If we can keep in mind these facts about what is asserted in a disjunction, and about the conditions under which a disjunction is true, then reasoning about alternatives is pretty straightforward.

5.1.2 Denying a Disjunct

Let's study this kind of reasoning by starting with a simple example. When Sherlock Holmes tries to figure out who the murderer is, he starts by listing the possibilities and then ruling them out one by one until only one possibility remains. His reasoning is something like this:

I know that the maid, the butler, or the chauffeur is the murderer. But I now have good reason to believe that the maid was in town buying meat at the time of the murder and that the chauffeur was driving to the train station to pick up the general's daughter. This leaves the butler as the only remaining possibility. So he must be the murderer.

If we wanted to formulate this reasoning in abstract terms, we could rewrite it as follows:

Either P, or Q, or R is the case.
But it is not the case the P.
And it is not the case that R.
So, it must be the case that Q.

Because the second and third premise involve ruling out a possibility, we can call this form of reasoning about alternatives **denying a disjunct**. This kind of reasoning starts by listing a series of possibilities in the form of a disjunction, and then denying one or more of the disjuncts, concluding that the remaining disjunct must be true. Using the False Premise Test we studied in Chapter 3, we can see that these premises are dependent: they work together to support the conclusion. If any of the premises were false, the remaining ones would not provide any support for the conclusion. In reasoning about alternatives by denying a disjunct, the premises are always dependent. Using the tests we studied in Chapter 3, we can also see that denying a disjunct is a valid form of reasoning: if the premises are true, the conclusion has got to be true too. If Sherlock is right that one of those three did it, and if he is


right that neither the maid nor the chauffeur did it, then he must also be right in concluding that the butler did it. There is no other possibility. If the doctor is right that patient either has a viral infection or cancer, but is able to rule out cancer, then it must be that the patient has a viral infection. Denying a disjunct is thus a powerful way to reason about alternatives because it is relatively easy to use and is always valid.

BOX 5.1.2 SUMMARY: DENYING A DISJUNCT

Denying a disjunct is reasoning by ruling out a possibility and concluding that the remaining possibility must be the case. It is always valid.

It has the following symbolic form:

1. A or B
2. It is not the case that A
3. So, it must be the case that B



5.1.3 False Disjunctions

Unfortunately, there are several ways that reasoning about alternatives can go wrong. As we know, we always want more from our reasoning than just validity. Not only do we want our reasons to support our conclusion, which is what validity guarantees, we also want our reasons to be **acceptable**. In the case of reasoning about alternatives, this means that we also want our disjunctions and our claims denying a disjunct to be true or at least acceptable. An argument that has a false disjunction as a premise makes the mistake that we call the **Mistake of the False Disjunction** (sometimes also called the Mistake of the False Dilemma). The risk of a false disjunction poses a special problem for reasoning about alternatives. You may recall that, in general, a valid argument with false premises might still have a true conclusion; one can get the right answer even if one's evidence is inaccurate or misleading. But in the case of denying a disjunct, which is always valid, this general rule fails. If the disjunction is false—that is, if neither disjunct is true—then reasoning by denying the disjunct will inevitably yield a false conclusion. Here is an example:

The car won't start. Either its battery is dead or it is out of gas. I just checked and there is plenty of gas. So, the battery must be dead:

This reasoning is valid and the disjunction is the proposition that either the battery is dead or the car is out of gas. Let's suppose that this disjunction is false, that is to say that the battery is fine and the car has lots of gas. This means that the argument's conclusion—that the battery is dead—is also false. The reason the car won't start, let us suppose, is that the spark plugs have been removed. Reasoning by denying a disjunct will yield a false conclusion when the disjunction itself is false. Consider a more serious example. Suppose that the doctor is mistaken that one of those two conditions is the cause of the patient's symptoms. Perhaps the symptoms are caused by a bacterial infection. In that case, when she rules out cancer and concludes that it is a viral infection, her conclusion is mistaken, even though her reasoning was valid. So we want to make sure that our reasoning about alternatives does not involve the mistake of false alternative. We want to make sure that our disjunctions are true. Making sure that our disjunctive premises are true when denying a disjunct is important not just because having true premises is always important to having a good argument, but because if the disjunction is false, then the conclusion of our reasoning will be false too.

BOX 5.1.3 MISTAKE TO AVOID: FALSE ALTERNATIVE

It is a mistake for an argument that reasons about alternatives to have a false disjunction. Not only does this make the premise unacceptable, but it also makes the conclusion false. While such arguments are valid, their conclusion is always false.

5.1.4 When Are Disjunctions Acceptable?

One way to know whether a disjunction is true is to know which of its disjuncts is true. But if we knew this, then we would not need to reason about alternatives, since we would already know which alternative was the truth. If we already knew that the spark plugs had been removed, then we would know that the disjunction was false. But then there would be no need to reason about alternatives. The reason the doctor is considering different possibilities is that she does not know which of the several possible causes the actual one is. One way to make sure that our disjunctions are true, even when we do not know disjunct is true, is to make sure that they are **exhaustive**. A disjunction is exhaustive when it includes all the possibilities that have not yet been ruled out. If we make sure that our reasoning about alternatives includes all the

possible alternatives, then we will be guaranteed to have a true disjunction. Exhaustive disjunctions help us to avoid the Mistake of the False Disjunction.

Making sure that we have an exhaustive disjunction can also help us to avoid a second and related mistake, the **Mistake of the Lucky Disjunction**. Suppose that the doctor ignored or overlooked other possible conditions that could cause the very same symptoms. Suppose those conditions could have been caused by a bacterial infection, by a bladder condition, and by a vitamin deficiency, and that she overlooked or ignored these possibilities. But suppose that as a matter of fact it is a viral infection causing the patient's symptoms. If we later found out that the doctor had overlooked or ignored these possibilities, then we would be very disappointed. We would feel that she ought to have been more careful to rule out all the relevant possibilities. She should have ordered more tests than she did. We would say that she did not really know that the patient had a viral infection; she just got lucky. She was lucky that the disjunction she relied on in her reasoning about the possibilities just happened to be true. She should have done a more exhaustive study; she should have relied on an exhaustive disjunction. So exhaustive disjunctions are preferable not just because they are guaranteed to be true (and so help us to avoid the Mistake of the False Disjunction), they are also preferable because they eliminate the need to rely on luck in our reasoning, and so help us to avoid the Mistake of the Lucky Disjunction.

But how can we be sure that a disjunction is exhaustive? How can we be sure that we know what the relevant possibilities are? This is very difficult. Sometimes, even the experts in a field are not sure just what they are. It can take years of study to discover all the possible causes of some set of symptoms. It can require a highly trained and flexible imagination—the kind that scientific geniuses like Galileo, Newton, and Einstein and artistic geniuses like Picasso and Pollock are reported to have had—to conceive of new possibilities, to think “outside the box.” Because it is so difficult to know when we have identified all the possibilities, part of our duty as critical thinkers is to keep an open mind about what is possible. More specifically, we have to keep in mind that just because we do not know of or even cannot conceive of other possibilities this does not mean that there are no other possibilities. Believing that there are no other possibilities on the grounds that one does not know or cannot think of any more is to commit a special version of the **Mistake of Appealing to Ignorance**: it is a mistake to believe that something is the case simply because it has not been proven otherwise. This mistake can occur in reasoning about any topic

at all, but is especially common in reasoning about alternatives. Here is an example:

James must have spilled the milk. Either he did it or the cat did it. What other possibility could there be? And there is no way the cat did it, since she would be covered in milk if she had. So this is why I am sure that James spilled it.

It might be that those are the only two possibilities. But the fact that the author could not think of others is not sufficient reason to believe that those are the only two possibilities. Making sure that our reasoning about alternatives involves exhaustive disjunctions can be exhausting—but that is no reason not to try. And it is certainly no reason to ignore the possibility that we have overlooked some possibilities.

BOX 5.14 MISTAKES TO AVOID: APPEALING TO IGNORANCE

Believing that there are no other possibilities on the grounds that one does not know or cannot think of any more is to commit a special version of the **Mistake of Appealing to Ignorance**: it is a mistake to believe that something is the case simply because it has not been proven otherwise.

5.1.5 Exclusive Disjunctions

We have been discussing the idea of an exhaustive disjunction. An exhaustive disjunction is one that includes all the possibilities that have not been ruled out. A disjunction can also be **exclusive**. A disjunction is exclusive just in case at most one of its disjuncts is true. The following are examples:

Jane's baby is a boy or a girl.

The soccer team won, tied, or lost.

The symptoms are either caused by cancer or by something else.

When it comes to babies, they are either boys or girls, and never both. (OK. Maybe this is not the best example, given the fact that there are hermaphrodites; but you get the idea.) Sometimes when we know that a set of possibilities is exclusive we still use a disjunction to state the

alternatives. But this is a bit misleading, since in asserting a disjunction we are sometimes allowing that more than one of the disjuncts is true. If we know that at most one is true, we should say this explicitly. Usually we don't have to because we know that our listeners also know that the disjunction is an exclusive one. But it is even more important to make this explicit when we reason using disjunctions that we know are exclusive. To see why, consider the following argument:

Either the maid did it or the butler did it. The butler just confessed. So the maid is innocent.

In this argument, instead of denying a disjunct, Sherlock asserts that one of the disjuncts is true and then concludes that the other one is not true. This form of reasoning is called **affirming a disjunct**. Unlike denying a disjunct, which is always valid, affirming a disjunct is valid only when the disjunction is exclusive. Only if one and only one of the disjuncts is true does it follow from the fact that one is true that the other is false. We can rule out the other possibilities only when we know that the possibilities are incompatible. To make matters even more complicated, it is possible for a disjunction to be partly exclusive: only some of the disjuncts are incompatible one with another. For example:

To balance the books, the government has several possibilities: it can raise income tax rates; or cut spending; or borrow money; or it can cut income tax rates and hope that that stimulates the economy enough to raise the needed revenues.

A government could both raise taxes and cut spending in order to balance the books, but it cannot raise and lower the income taxes at the same time. So a disjunction can be partly exclusive or wholly exclusive. This makes reasoning with them more complicated than denying a disjunct. What is more, just as it is often hard to know whether a disjunction is exhaustive, it is also hard to know whether a disjunction is exclusive. This often takes specialized knowledge. Moreover, if we know that the disjunction we are reasoning with is exclusive, then this bit of knowledge is in effect a premise that we are relying on when we reason by affirming a disjunct, and we should make this premise an explicit part of our reasoning. If Sherlock knows that the murderer acted alone, then he should make this clear.

Either the maid did it or the butler did it. And I know that the murderer worked alone. And the butler just confessed. So, the maid is innocent, after all.

In this bit of reasoning, the premises are dependent just as with denying a disjunct. And it is clear that here the reasoning is valid. If we were to write it out in symbolic form, it would look something like this:

Either P or Q is the case.
 But only one of them can be the case.
 P is the case.
 So Q is not the case.

This form of reasoning is always valid, just as is denying a disjunct, which is a good thing. Even better, the argument makes explicit all the evidence the author was relying on in drawing the conclusion. Reasoning using disjunctions that we know are exclusive is more complicated than reasoning about alternatives by denying a disjunct. This does not mean that we should avoid affirming a disjunct, but only that when we do we should make this explicit so that we can make it a question whether we are right that the disjunction really is exclusive.

BOX 5.1.5 MISTAKES TO AVOID: AFFIRMING A DISJUNCT

Affirming a disjunct is reasoning from the truth of one disjunct to the conclusion that the other disjunct(s) is false. It is valid only if the disjunction is an exclusive disjunction. But if one knows that the disjunction is exclusive, then one should add this piece of information as an additional premise.

5.1.6 Criticizing Reasoning about Alternatives

We have been discussing reasoning about alternatives, and we have seen that there are several ways that this kind of reasoning can go wrong. When we reason on our own about alternatives we ought to keep these errors in mind. We also need to keep them in mind when we assess someone else's reasoning. If we read an argument in the newspaper or in a book where the author reasons about some alternatives we should ask ourselves whether the author is committing any of the mistakes we have identified here. Is her disjunction true? Is it exhaustive? Is she claiming the disjunction is true just because she cannot think of any other possibilities? Is she affirming a disjunct even though her disjunction is not exclusive? If we decide that the author has committed one or more of these mistakes, then we have the obliga-

tion to say exactly which mistake it is and we have to be ready to back this up with evidence. If we claim that someone has ignored or overlooked some possibilities, and so charge them with using a disjunction that is not exhaustive, then we have to be ready to make good on this claim either by pointing out some specific possibilities that have been overlooked or ignored or else provide some other reason to think the disjunction is not exhaustive. We also have to be careful not to ridicule the author's view by noting possibilities that are extremely unlikely or unreasonable. Usually, there are many possibilities that are simply out of the question—so far from what might actually happen—that they are ruled out without argument. They are not, we might say, “live options.” Perhaps building a land bridge across the river might be an alternative to either building a new bridge or renovating the existing one. But, depending on the discussion, it might be that it is simply not a live possibility. In that case, it would be unreasonable—a form of the red herring—to charge an author with relying on a nonexhaustive disjunction if this was the only possibility one could mention. This would be a case of the **Red Herring Fallacy** since one would be introducing an irrelevant possibility simply in order to better criticize the author's position.

BOX 5.1.6 DECIDING WHAT TO DO: IDENTIFY ALTERNATIVES

Thinking critically about what to do invariably requires thinking about alternatives. For usually one has to decide from among several competing ends, and there are almost always several different means for achieving any given end. As a general rule of thumb, one has not thought carefully enough about what to do unless one has identified three possible courses of action.

EXERCISE 5.1

1. Short-answer questions:

- a. What does a disjunction assert?
- b. What is the difference between an exclusive and a nonexclusive disjunction? Construct three examples of each kind.
- c. Why must the conclusion of an argument that denies a disjunct be false if the disjunction is false too? Use examples to illustrate your answer.

- d. When is affirming a disjunct valid? Give an example.
 - e. What is the symbolic form of denying a disjunct?
 - f. Are the premises in denying a disjunct dependent or independent? Explain, using an example.
 - g. What is an appeal to ignorance, and why is it a critical thinking mistake?
 - h. What is the red herring fallacy, and why is it a critical thinking mistake?
 - i. What is an exhaustive disjunction? Give two examples.
2. Construct five disjunctions that are true and exclusive but that a reasonable person would think are not exhaustive.
3. In the following arguments, identify the premises and conclusions and then determine whether the argument is valid. If it is not valid, identify the error.
- a. The city can either build a new bridge or else renovate the new one. But building a new bridge is so expensive that it cannot be afforded. We will have to live with a renovated one.
 - b. A father to his child: The restaurant's menu says that you can have cake or apple pie for desert. You have already ordered the cake, so you cannot have the pie too.
 - c. My boss said that I had to do the dishes or the laundry, neither of which I wanted to do. I just finished the dishes, so I do not have to do the laundry.
 - d. Either humans evolved from other species or else God created us. I believe that God created all life. So humans did not evolve from other species.
 - e. I cannot walk and chew gum at the same time, and I am walking now. So I had better not chew gum.
 - f. To fight global warming, we can either raise taxes on fossil fuels or else find alternative sources of energy right now. Since there is no chance of finding alternative sources of energy now, we must raise taxes.
 - g. According to the computer models, either the economy will grow by 2% in the next quarter or else unemployment will continue to rise. Since unemployment cannot rise any higher than it is, the economy is probably going to grow.
4. Look through the letters to the editor in your local newspaper or magazine and find five examples of reasoning about alternatives.

Analyze them by identifying their premises and conclusions. In your judgment, are they examples of good reasoning? Do they commit any of the mistakes that we have discussed?

5.2 REASONING ABOUT NECESSARY AND SUFFICIENT CONDITIONS

Sometimes when we reason about what to believe or what to do, we reason about necessary and sufficient conditions. Here are some examples:

If the city decides to build a new bridge, then it will have to raise taxes to pay for the construction, and raising taxes will negatively affect the business climate, so the city should not build a new bridge.

If the patient's symptoms had been accompanied by a fever that had lasted for more than three days, then it would have been caused by a bacterium rather than by a virus, but it was not. So the illness is probably viral.

The restaurant claims to provide excellent service, but this requires careful attention to detail, which the waiters completely lack, so it is simply not true that the restaurant provides excellent service.

I want to keep my grass green, and I learned from the gardening book that grass needs regular watering to thrive. So I plan to water it regularly.

I need to pass this course, and studying hard for it is the only way to guarantee that I'll pass. So my plan is to study really hard.

If we leave Iraq now, a civil war will break out and the neighboring countries will probably enter in, leading to a vast and bloody conflict far worse than what will happen if we stay. So we must stay come what may.

These are all cases of reasoning about what is necessary for what or about what is sufficient for what. Their range shows that this kind of reasoning occurs in a wide variety of disciplines. And as these cases illustrate, we commonly reason about necessary and sufficient conditions using conditional sentences, sentences with an "if" and a "then." In this section, we will study the nature of this kind of reasoning and identify some of the mistakes that can arise.

5.2.1 The Meaning of Conditionals

We saw that we can make our reasoning about alternatives more reflective if we keep in mind what we learned in Chapter 2 about the meaning of disjunctions. The same is true for reasoning about necessary and sufficient conditions. If we keep in mind what a conditional sentence is used to say, and under what conditions that kind of sentence says something true, then this will make it easier to understand reasoning about necessary and sufficient conditions and to identify, and so avoid, some of the possible mistakes. So let's start by reviewing some of what we know about the meaning of "if, then" assertions.

Here are some examples of conditionals:

If Stephen is the prime minister of Canada, then Stephen is a politician.

If you want the lawn to be green, then you have to water it.

If a restaurant has excellent service, then its waiters pay careful attention to every detail.

The part of the conditional that follows the "if" we call the **antecedent** and the sentence that follows the "then" we call the **consequent**. In making an assertion using a conditional, we do not assert either the antecedent or the consequent. That is, someone who claims that if Stephen Harper is prime minister of Canada, then Stephen Harper is a politician is not thereby claiming that Stephen Harper *is* the prime minister of Canada or that Stephen Harper *is* a politician. She might well believe these things, and she may even be willing to assert them. But she is not asserting them in using that conditional to make an assertion. Instead, she is claiming that there is a relation of a certain kind between the antecedent and the consequent. More specifically, she is claiming that if the antecedent is true, then the consequent is true too. It is helpful to state this relation using the notions of necessary and sufficient conditions. In asserting a conditional, one is asserting two things. First, one is asserting that the truth of the antecedent is **sufficient** for the truth of the consequent. Second, one is asserting that the truth of the consequent is **necessary** for the truth of the antecedent. Let's look at each of these points in turn.

5.2.1.1 Sufficient Conditions First, if a conditional is true, then the truth of the antecedent is sufficient for the truth of the consequent. We are already familiar with the logical idea of sufficiency. It is the idea behind the notion of validity. To say that the truth of the antecedent is

sufficient for the truth of the consequent just means that all that it would take for the consequent to be true is for the antecedent to be true. One way for Stephen Harper to be a politician is for him to be prime minister of Canada. Being prime minister is sufficient for, or guarantees, that one is a politician. One way to become a millionaire is to win the lottery. Winning the lottery is sufficient for becoming a millionaire. If you win the lottery, you will become a millionaire. One way to pass this course is to get an A. If you get an A, then you will pass. Getting an A is sufficient for passing this course.

These examples also can help us to see another important point. Something can be sufficient for something without being necessary for it. Winning the lottery is one way to become a millionaire; but it is not the only way. Inheriting from a rich uncle is just as good. Working hard and saving money is another way to become a millionaire. Robbing a casino may make one a millionaire (if only for a few moments). Any of these conditions is sufficient for becoming a millionaire, but none is necessary. Likewise, getting a C in this class—or even a D—is also sufficient for passing it. There are often, as the saying goes, many ways to skin a cat.

5.2.1.2 Necessary Conditions The second point is that if a conditional is true, then the truth of the consequent is necessary for the truth of the antecedent. In other words, if the conditional is true, then there is no way that the antecedent could be true if the consequent were false. If you are not a millionaire, then you did not win the lottery. If you did not pass this course, then you did not get an A. Just as there can be many sufficient conditions for something, so there can be many necessary conditions for it too. Having a healthy lawn requires more than just lots of water; it also requires lots of sunshine, fertilizer, and an absence of grubs. Becoming a millionaire requires having more than 100 dollars, and it also requires having more than 200 dollars. It also requires a system of property law to establish ownership and a stable currency system to make money possible. To bake cookies you need more than just butter and flour, you also need some sugar and vanilla and chocolate chips (these are needed if you want the cookies to be any good, anyway). Usually, there are lots and lots of necessary conditions for something.

5.2.1.3 Necessary and Sufficient Conditions A condition that is necessary for something can also be sufficient for it. Jones' being an unmarried male is necessary for his being a bachelor, but it is also sufficient. To be a bachelor, you must be a male and unmarried. Those

two conditions are each necessary, though neither one is sufficient on its own for being a bachelor. But together, they are sufficient. We can put this by saying that those conditions are individually necessary and jointly sufficient. Sugar, butter, and flour are each necessary for having cookie dough, and they are jointly sufficient. All you need to have some cookie dough is sugar, flour and butter. We can state the fact that certain conditions are individually necessary and jointly sufficient using a conjunction of conditionals:

If Jones is a bachelor, then he is an unmarried man *and* if he is an unmarried man, then he is a bachelor.

But it is usually easier just to join the conditionals with what we call a “bi-conditional,” as follows:

Jones is a bachelor **if and only if** Jones is an unmarried man.

We use bi-conditionals to say that certain conditions are both necessary and sufficient for something.

As is already clear from our discussion, there are many ways in English to state a necessary or sufficient condition. It is worth spending a bit of time noting some of them. We can state a sufficient condition using the “If ... , then ...” form of a conditional as follows:

If Stephen is prime minister of Canada, then Stephen is a politician.

If Jones is a bachelor, then is Jones is a man.

If the lawn will grow, then it has been watered.

If you are a millionaire, then you have more than 200 dollars.

In these sentences, the focus is on stating a sufficient condition for the consequent. We can also put the focus on the necessary condition using the “only if” form of conditional, as follows:

Stephen is prime minister of Canada only if Stephen is a politician.

Jones is a bachelor only is Jones is a man.

The lawn will grow only if it is watered.

You are a millionaire only if you have more than 200 dollars.

Because of the intimate relation we have noted between necessary and sufficient conditions, these are (for most intents and purposes, anyway)

just two ways of saying the very same thing. In general, the following are just two ways of saying the same thing:

If P, then Q

P only if Q

Sometimes, we reverse the order of the two parts of the conditional, as follows:

I will go to the party, if Jones goes too.

If Jones goes to the party, then I will go too.

Only if Jones is a man can Jones be a bachelor.

Jones is bachelor only if Jones is a man.

The first and second sentences identify Jones' attendance at the party as a sufficient condition for the speaker's going to the party. The third and fourth sentences identify Jones' being a man as a necessary condition for his being a bachelor. The word "unless" is also used to state a necessary condition.

Jones is married, unless he is a bachelor.

If Jones is married, then he is not a bachelor.

If Jones is a bachelor, then he is not married.

These are just three ways of saying the same thing: that being unmarried is a necessary condition for being a bachelor. The important thing to keep in mind as one reads a conditional sentence that is not in the nice and neat "if, then" form is to ask what is being said to be necessary or sufficient for what. One strategy is to begin by identifying the two conditions, and then asking which the author is saying is sufficient for which. This is sometimes difficult to do if, as does happen, the antecedent and consequent in the conditional are not complete sentences. Sometimes, the author leaves out some of the implicit information, just for ease and because her audience can figure it out. When analyzing a bit of reasoning, it is always best to rewrite the conditional inserting that missing information.

EXERCISE 5.2.1

1. Short-answer questions:

- a. What does a conditional assert? Use the four-step definition method discussed in Chapter 2 to answer this question, and use examples to illustrate your answer.

- b. Could a conditional be true if its consequent is false? Use an example to illustrate.
 - c. What is the difference between a necessary and a sufficient condition? Use the four-step definition method discussed in Chapter 2 to make your answer clear. Use examples.
 - d. Give an example of something that is necessary but not sufficient for something else.
 - e. Give an example of something that is both necessary and sufficient for something else.
 - f. Give an example of something that is a necessary part of a sufficient but not necessary condition for something.
2. For each of the following, find five things that are necessary but not sufficient for it:
 - a. Being a father
 - b. Being happy
 - c. Being a good movie
 - d. Murder
 - e. Love
 - f. Being president
 - g. Passing this class
3. For each of the following, find something for which it is sufficient but not necessary:
 - a. Having 25 dollars.
 - b. Getting an A in this course.
 - c. Being president of the United States.
 - d. Being the spouse of a supermodel.
4. For the following conditionals, find a conditional of the “If, then” form that means the very same thing.
 - a. Jones will go to work only if Jones gets paid today.
 - b. I would have retired if someone had asked me to.
 - c. The economy will rebound only if taxes are cut.
 - d. The economy will not rebound unless taxes are cut.
 - e. The economy will not rebound if taxes are not cut.
 - f. The movie will not be entertaining unless it is a musical.
 - g. The movie will be entertaining unless it is a musical.
 - h. The grass won’t grow unless you water it.
 - i. To make grass grow you need to water it.
 - j. Eating lots of vegetables is essential to a healthy diet.

- k. To be good at critical thinking you need to trust but verify.
- l. The burger is adequately cooked when its internal temperature is 150 degrees.
5. In the following arguments, identify the premises and conclusion. Rewrite any conditionals in an “If, then” form that means the same thing. If a conditional premise is implicit (or missing) insert it.
- If someone had asked me to resign I would have. But no one ever asked me to, so that is why I did not.
 - If I had gone to the party, Jones would have seen me. If he had seen me, he would have been pretty angry. So I didn’t go.
 - The grass is growing very well. I guess that someone watered it.
 - If humans were alive when the dinosaurs roamed the earth, there would be evidence of this in the fossil record. But there is no such evidence. So I think humans were not alive then.
 - Success requires hard work, and I’m working hard, so I’ll succeed.
 - The recipe says that the cake is done when it starts to pull away from the sides of the pan, and this is exactly what is happening, so the cake is done.
 - The chain on my bike is very rusty, so it is about to break.
 - Vitamin C helps to avoid colds, And Emily is very sick, so she has not been taking vitamin C.

5.2.2 Valid Forms of Reasoning about Necessary and Sufficient Conditions

Now that we have studied necessary and sufficient conditions and how we can use conditional sentences to formulate them, we can study how to reason about them. Just as with reasoning about alternatives, reasoning about necessary and sufficient conditions is very common and relatively straightforward. However, just as with reasoning about alternatives, there are some ways to go wrong that we should be aware of. Once we have studied this kind of reasoning at an abstract level, we will consider two special applications of this kind of reasoning. First, we will study reasoning involving standards and definitions. Then we will look at reasoning about causes.

Let’s begin our discussion with a great philosophical chestnut:

Aristotle is man. If Aristotle is a man, then Aristotle is mortal. So, Aristotle is mortal.

As always, before we assess an argument we should analyze it into premises and conclusion. In this case, the presence of the indicator

word “so” tells us that the final sentence is the conclusion. The first two sentences are thus premises. Using the false premise test, we can see that in this case the premises work together; they are dependent. If either were false, the other would provide no support at all for the conclusion. Suppose that it were not true that Aristotle is human. In that case, the fact that that if he was human then he would be mortal would provide no reason at all to think that Aristotle is mortal. Likewise, suppose that being human is not sufficient for being mortal. In that case, the fact that Aristotle is human would provide no reason to think that he is mortal. So the premises are supposed to work together.

Now that we have analyzed the argument, we can assess it. As always, there are two questions to ask: do the premises support the conclusion and are they acceptable? Let’s start by considering whether the premises support the conclusion. If we keep in mind what we learned about necessary and sufficient conditions, then figuring out whether the premises support the conclusion in a case of reasoning about them is usually pretty straightforward. Part of what makes this age-old example of an argument a wonderful philosophical chestnut is that it is pretty obviously a valid argument. There is no way that its premises could be true and its conclusion false. If he really is a man, and if it really is true that if he is a man then he is mortal, then there is no way for him not to be mortal. The truth of the premises guarantees that of the conclusion. If you think back to what we said about the meaning of conditionals in the previous section, it should not be at all surprising that this argument is valid. After all, the conditional that is the second premise in the argument simply says that Aristotle’s being a human is sufficient for his being mortal. And the first premise simply states that that sufficient condition for his being mortal in fact does obtain. If his being a human is enough to make him mortal, and if he is human, then how could he not be mortal? So the argument is valid.

BOX 5.2.2A SUMMARY: MODUS PONENS

Modus ponens is a valid form of reasoning about necessary and sufficient conditions. It has the following symbolic form:

- | | |
|-------------------|-----------|
| (1) If P, then Q. | 1 + 2 |
| (2) P | ↓ |
| So, (3) Q | 3 |

This form of reasoning about necessary and sufficient conditions is called **modus ponens**, or affirming the antecedent. Any argument of that form is valid. The validity of modus ponens turns on the fact that if a conditional is true then the antecedent is sufficient for the consequent.

Another valid form of reasoning about necessary and sufficient conditions turns on the fact that if a conditional is true, then the truth of the consequent is necessary for that of the antecedent. Any argument of the following form is also valid. Here is an example:

If Jones is a bachelor, then he is unmarried.
 But Jones is married.
 So, Jones is not a bachelor.

This form of reasoning is called **modus tollens**, or denying the consequent. It is valid too. If it is true that being unmarried is necessary for being a bachelor, which is what the conditional premise asserts, and if it is true that Jones is married, which is what the other premise asserts, then it has to be true that Jones is not a bachelor. This form of reasoning is valid just because if a conditional is true, then the consequent must be true if the antecedent is.

BOX 5.2.2B SUMMARY: MODUS TOLLENS

Modus tollens is reasoning from the absence of a necessary condition to the absence of a sufficient condition. It has the following symbolic form:

- (1) If P, then Q
 (2) It is not the case that Q.
 So, (3) it is not the case that P.



The two forms of reasoning we have look at so far involve claims about necessary and sufficient conditions as premises, but we can also reason to conclusions about necessary and sufficient conditions. Here is a valid form of this kind of reasoning:

If Aristotle is human, then Aristotle is mortal.
 If Aristotle is mortal, then Aristotle is not God.
 So, if Aristotle is human, then Aristotle is not God.

In this form of reasoning, which we can call a **pure conditional** (we can call it that since all the premises and the conclusion are conditionals), we reason along from one pair of sufficient conditions to a third one. If being human is sufficient for being mortal, and if being mortal is sufficient for not being God, then it has to be true that being human is sufficient for not being God. So this form of reasoning is valid too.

BOX 5.2.2C SUMMARY: PURE CONDITIONAL REASONING

Pure conditional reasoning is reasoning from two or more conditionals to a conditional. It is valid if it is of the following form:

<p>(1) If P, then Q</p> <p>(2) If Q, then R</p> <p>(3) So, if P, then R</p>	<p>(1) + (2)</p> <p>↓</p> <p>(3)</p>	
---	--	--

We have seen three valid forms of reasoning about necessary and sufficient conditions. There are of course mixed cases that are valid too.

If we raise taxes to pay for the new bridge, then the public will turn against the construction. If the public turns against the project, then it will surely fail. But we cannot let it fail. So we should not raise taxes to pay for it.

This bit of reasoning incorporates Pure Conditional reasoning with Modus Tollens. We could put the reasoning symbolically as follows:

If P, then Q
 If Q, then R
 But it is not the case that R
 So, it is not the case that P.

This example of mixed reasoning using both modus tollens and a pure conditional is valid.

5.2.3 Invalid Forms of Reasoning about Necessary and Sufficient Conditions

We have identified three very common valid forms of reasoning about necessary and sufficient conditions. There are also some invalid forms that are worth noting that come from confusing a necessary condition for a sufficient one or a sufficient one for a necessary one. Here is one example:

If Jones has diabetes, then he shouldn't be eating that cake and ice cream.
But he doesn't have diabetes, so he can eat some of it.

On its face, this bit of reasoning, which we can call **denying the antecedent**, can seem very persuasive. If you look through letters to the editor or listen to friends talk, you will probably find lots of examples of it. But this reasoning is invalid. To see this, think about it in terms of necessary and sufficient conditions. The first premise says that having diabetes is sufficient reason for avoiding eating sweet desserts. Suppose that is true. But that premise does not say that having diabetes is necessary for avoiding sweet desserts. It does not say that only people with diabetes should avoid desserts. Indeed, we know that there are lots of other good reasons for sometimes avoiding sweet desserts. Maybe Jones is on a strict diet. That would be good reason. Maybe Jones is lactose intolerant, and so should avoid ice cream. Maybe Jones' wife is avoiding sweets and by eating them himself he would weaken her resolve, which he does not want to do. There are lots of conditions that are sufficient for avoiding that kind of dessert, and Jones might meet one of those conditions but not another. So even if the premises are true, even if he does not have diabetes, the conclusion of this bit of reasoning might still be false. In other words, arguments of this form are not guaranteed to be valid in the way that arguments of the form of Modus Ponens are guaranteed to be valid. This mistake comes from taking a sufficient condition to be a necessary one.

Here is an example of another invalid form of reasoning about necessary and sufficient conditions, one that comes from taking a necessary condition to be a sufficient one:

Susan feels sick most mornings, which is exactly how she would feel if she was pregnant. So, I think she's going to have a baby.

Let's first analyze this into premises and conclusion. The presence of the conclusion indicator "so" tells us that the final sentence is the conclusion. The first complete sentence—that Susan feels sick most

mornings—is one premise. The other premise is stated in a complex way. But it can be reformulated as follows:

If Susan were pregnant, then she would feel sick most mornings.

Now we can see that the reasoning has the following form:

If P, then Q
Q
So, P

This form of reasoning is called **affirming the consequent**. But it is not valid. The first premise—the conditional one—says that Q is a necessary condition for P. That means that for P to occur or be true, Q has to occur or be true too. But we know that a necessary condition does not have to be a sufficient one. So, even if Q is the case, the fact that it is necessary for P does not show that P also obtains. This reasoning seems persuasive only if we confuse a necessary condition for a sufficient one. Suppose that it is true that if Susan is pregnant then she will feel sick in the morning. And suppose that she does feel sick in the morning. Couldn't it just be a coincidence that she is sick? Maybe she is suffering from some other condition that makes her sick in the morning. Even if feeling sick in the morning is necessary for being pregnant, it need not be sufficient. So this reasoning is not valid either.

BOX 5.2.3 MISTAKES TO AVOID

Denying the antecedent and affirming a consequent are not valid forms of reasoning.

Denying the antecedent	Affirming a consequent
If P, then Q	If P, then Q
Not P	Q
So, not Q	So, P

We have identified some valid forms of reasoning about necessary and sufficient conditions. We use Modus Ponens when we know that a condition that is sufficient for something is true. We use Modus Tollens when we know that a condition that is necessary for something is

absent. We use Pure Conditional to trace links between a series of sufficient conditions. So long as we can keep clearly in mind the difference between a necessary and a sufficient condition and the way we use our words to state and formulate those kinds of conditions, our reasoning about them we can avoid the mistakes and make sure that our conditional reasoning involves premises that support the conclusions we draw.

5.2.4 Making It Explicit

All of the examples we have looked at so far have conditional sentences that make explicit the author's (sometimes mistaken) assumptions about necessary and sufficient conditions. But we know that there are other ways to make claims about such conditions. It is helpful when analyzing such an argument to re-write the claim in conditional form. Here is an example:

To get a good job Jones will need to go to university. But he won't get accepted in any university. So he's not going to get a good job.

The conclusion of this argument is that Jones will not get a good job. None of the premises is formulated as a conditional sentence, but the first sentence does make a claim about a necessary condition. It claims that going to university is necessary for Jones to get a good job. We could formulate it in a conditional in either of the following ways:

Jones will get a good job only if he goes to university.

If Jones gets a good job, then he will have gone to university.

Now that we have rewritten that premise, we can rewrite the entire argument as follows:

Jones will get a good job only if he goes to university. But he will not go to university, so he will not get a good job.

Now we can more easily see that this argument has the form of Modus Tollens, and so is valid. Depending on whether the premises are true, this might be a good argument for the conclusion. As this example illustrates, it can be helpful in assessing reasoning about necessary and sufficient conditions if we can re-write the relevant claims in the form of conditionals.

Here is another example:

The restaurant does not deserve five stars. The food was too greasy.

Here the author does not make any claim at all, at least not explicitly, about necessary or sufficient conditions. But it is plausible to think that the author is assuming that the fact that the restaurant's food is so greasy disqualifies it from deserving five stars. In other words, the author is assuming that not having greasy food is necessary for deserving five stars. We can formulate this assumption as follows:

A restaurant deserves five stars only if its food is not greasy.

We could also formulate it as follows:

If a restaurant's food is greasy, then it does not deserve five stars.

Now we can see that the author's reasoning, with the missing premise added, can be formulated in the following way:

If a restaurant's food is greasy, then it does not deserve five stars. That restaurant's food is greasy. So, it does not deserve five stars.

Again, now we can see that the reasoning is valid. Whether the argument is ultimately good depends on whether the two premises—the one about what is required to deserve five stars, and the other about the food's greasiness—are true (or at least acceptable).

EXERCISE 5.2.4

1. Short-answer questions:

- a. What is modus ponens?
- b. Using the concepts of a necessary and a sufficient condition, explain:
 - i. Why modus ponens is valid. Give two examples.
 - ii. Why modus tollens is valid. Give two examples.
 - iii. Why denying the antecedent is not valid. Give two examples.
 - iv. Why affirming the consequent is not valid. Give two examples.

2. Look back at the arguments in Exercise 5.2.1. Assess them for validity.
3. Look at letters to the editor and find five letters that involve reasoning about necessary and sufficient conditions. Rewrite the reasoning into arguments involving conditionals, and assess them for validity.

5.2.5 When Are Claims about Necessary and Sufficient Conditions Acceptable?

As we know, we want more from our reasoning than just for the premises to support the conclusion; we want our premises to be true or at least acceptable. In the case of reasoning about necessary and sufficient conditions, we want our claims about these conditions to be true. Usually, this means that we want our conditional sentences to be true or at least acceptable. We know, in a general way, what it is for a conditional sentence to be true: it is true just in case the truth of the antecedent really is sufficient for the truth of the consequent; or, put the other way around, a conditional is true just in case the truth of the consequent really is necessary for the truth of the antecedent.

The **Counter-Example Strategy** can help us to decide whether a conditional is acceptable or true. A counter-example to a claim is an example that shows that that claim is false. A counter-example to a conditional would be an example that shows that the antecedent could be true even if the consequent were false. If we can imagine or conceive of a way that the antecedent might be true and the consequent false (or if we know for a fact that the antecedent is true and the consequent false), then we will know that the conditional itself is false. Imagining or conceiving of such a way is a matter of thinking of some alternative possibilities, ones where the world is the way the antecedent claims it is but not the way the consequent claims it to be. It is a bit like creating a little story of the world. Suppose that our conditional is this:

If Jones is a bachelor, then Jones is happy.

We might know that this is false because we know that Jones is a miserable bachelor who longs to marry Susan, who has spurned him more than once. In this case, we know both that the antecedent is true and that the consequent is false—this is enough to know that the conditional is false. But what if we do not know any of this? Suppose that we do not even know Jones very well. We can still assess the acceptability of the conditional by trying to imagine or conceive of the antecedent being true and the consequent false. In effect, this involves imagining

or conceiving of an alternative possibility, a kind of story of the world where Jones is a bachelor but he is unhappy. Doing this requires the same skills and strategies we discussed when we studied reasoning about alternatives. If we can come up with such a story, if we can identify a real possibility where the antecedent is true and the consequent is false, then we will have shown that the conditional is not acceptable. In this case, it is not that hard to do—after all, we all know lots of miserable bachelors. In other cases, it can be quite difficult to come up with such a counter-example. In this section we will look in some detail at two kinds of necessary and sufficient conditions—those that derive from definitions and those that involve causal relations—and we will study some strategies for deciding whether they are acceptable.

BOX 5.2.5 PRACTICAL TIP: COUNTER-EXAMPLE METHOD

An important step in assessing the acceptability of conditionals is looking for counter-examples. A counter-example is a case, either a real one or a fictional one, that shows that the conditional is false. A counter-example would be a case where the antecedent is true but the consequent is not (showing that the antecedent is not sufficient for the consequent). If you think that a conditional is false, then you need to present a counter-example and argue that it shows that it is.

But just one reminder: finding a counter-example to a claim proves that the claim is false; failing to find one does not prove that it is true. Suppose that we are unable to think of a counter-example to some conditional. This would not by itself prove that the conditional is true. Thinking that it would is a case of **appealing to ignorance**, which as we saw in our discussion of reasoning with alternatives is a mistake. The fact that we cannot think of a counter-example to a conditional is not a good reason to think it is true, just as our inability to think of a third alternative is not a good reason to think that a disjunction is true. Maybe our imaginations are just too limited.

5.3 REASONING WITH DEFINITIONS AND STANDARDS

We sometimes reason about necessary or sufficient conditions that derive from **definitions** in a fairly broad sense of that word. These

include ordinary dictionary definitions but also such things as standards for evaluation, systems for classification and categorization, legal definitions and the clauses and stipulations in contracts and agreements, standards for medical and other kinds of treatment, and rules for games and ordinary practices. Here are some examples, starting with one we have already seen:

A person is a bachelor only if he is unmarried.

A geometrical figure is a triangle if and only if the sum of its internal angles is 180 degrees.

A restaurant deserves five stars only if the service is exceptional.

If the fever persists for more than 24 hours, then call the doctor.

If players are on first, second, and third bases, then the bases are loaded.

If your filing status is Single, and at the end of 2006 you are under 65, then file an income tax return if your gross income was at least \$8,450, unless you are being claimed as a dependent on someone else's tax return.

As these examples illustrate, there is quite a range of different kinds of claims in this group, and assessing the acceptability of a conditional depends on what kind of claim it is. The acceptability of claims about rules and conventions requires looking at what the relevant rulebooks or contracts say. Appealing to a dictionary might settle whether some definition is acceptable. For technical terms in science or the law, appeal to experts might be needed. But even though different procedures are needed for deciding whether one kind of conditional or another is true or acceptable, when we reason with them we can follow the same strategies we have been discussing until now. Modus ponens, modus tollens, and the pure conditional are valid forms of reasoning for every kind of conditional.

Claims about what our words mean or what certain standards should be are simply claims about necessary and sufficient conditions. In Chapter 2 we discussed definitions, and we learned how to clarify and assess them. We learned that it is helpful to use the four-step definitional method to clarify what others and we mean by our words and claims. When we reason using definitions and standards, or when we analyze and evaluate someone else's reasoning of this kind, we can use the four-step method to clarify the definitions and standards. We saw that to assess whether a definition is acceptable we can use the counter-example method. Now we can see that these same methods

can be used to clarify, analyze, and assess reasoning using definitions and standards.

Reasoning about necessary and sufficient conditions that derive from definitions and standards runs the risk of **equivocation**. To equivocate is to use a word to mean different things. This is not always a bad thing, since our words often are ambiguous in a harmless and well-understood way. But there are two kinds of cases where it is a mistake. First, it is a mistake to equivocate in one's reasoning if one's words must mean one thing for the premises to be acceptable but another for them to support the conclusion. Here is an example:

Only man is a rational animal.
 Susan is not a man.
 So, Susan is not a rational animal.

The argument is valid only if the word "man" means the same thing in both premises. But the premises are acceptable only if it means something different. For the first premise to be true, it has to mean something like "human"; for the second premise to be true it has to mean "male human." This equivocation makes this into a bad argument.

Equivocation can also occur during discussions or in debates, and can make what are really disagreements over definitions and standards look like disagreements over the facts. Two movie reviewers who watch the very same film might disagree over its quality as a thriller. It might be that one knows something more than the other about the film, and that their disagreement derives from this difference. But it might instead be that they simply have different standards or ideas about what it means for a movie to be a "good thriller." Perhaps one thinks that the espionage film noir *The Thirty Nine Steps*, with its dark scenes and shady characters, is the standard against which to assess and compare thrillers, while the other thinks that *Vertigo*, with Jimmy Stewart hanging from the top of the tower, sets the standard. If that is the case, and if they do not realize that they have different standards in mind, then they are really talking past one another when they argue about whether a certain film is a good thriller or not. They mean different things, or have different standards in mind. There is nothing wrong with that, and it can become itself a topic of discussion, with one side trying to persuade the other to adopt its standard. But it would be a mistake not to recognize that the discussion involves an equivocation about standards.

BOX 5.3 MISTAKES TO AVOID: EQUIVOCATION

To equivocate is to use a word to mean different things. This is not always a bad thing, since our words often are ambiguous in a harmless and well-understood way. But there are two kinds of cases where it is a mistake. First, it is a mistake to equivocate in one's reasoning if one's words must mean one thing for the premises to be acceptable but another for them to support the conclusion. Second, it is a mistake not to recognize that participants in a disagreement mean different things by key words in their dispute.

The risk of unnoticed equivocation is also high in discussions over values and morality. Debates over the morality of abortion, for instance, sometimes involve disagreements about how to balance off the rights of the unborn with those of the pregnant woman and the father. Sometimes the debate turns on complex medical and physiological facts not well known by either side. But sometimes, debates over abortion seem to turn on hidden disagreements over what is to count as a person. Is a three-month-old fetus a person? In part this is a physiological question, since the physiological features of a three-month-old fetus are relevant. But it is not a question that can be settled through appeal to physiology. Even if we settled all of the physiological facts, we could imagine sincere and reflective people continuing to disagree over whether a three-month fetus "counts" as a person. This would be a case where the debate turns on a difference over definitions or standards, which is fine, so long as all sides recognize this for what it is.

What can be done to make progress when a discussion or debate turns on a difference over standards or definitions? It is important, first of all, to make the difference as clear as possible, and here the definitional method discussed in Chapter 2 can be useful. Formulating a clear statement of the definition or standard might, of itself, move the discussion along, as the two sides come to better understand what the other has in mind. Examples and contrasts can be useful too, though they can also sometimes be just as controversial as the original disagreement. But the goal should be to find as much common ground as possible. If the different sides in the discussion can get clear on where they agree—on the physiological issues, for instance, and on some clear examples of persons and of what makes them persons—then they can move on to discuss the controversial cases. It might be, though, that at bottom the discussion turns on competing visions of what is valuable, all sides

agreeing on the nature of something but disagreeing about its relative value. Such disagreements are among the hardest of all.

BOX 5.3B PRACTICAL TIP: LOOK FOR COMMON GROUND

In a discussion, it is always very important to find common ground, so that areas of disagreement can be clearly identified and resolved. Two kinds of common ground are important. First, one should find common factual ground: it is important in a disagreement to find what all sides agree on, in order to help them focus on their actual disagreements. But it is just as important to find common linguistic ground: it is important that all sides agree on how to use their words to say what the facts are or might be. To avoid the risk of equivocation in a discussion or debate, use the four-step definitional method from Chapter 2 to construct and evaluate definitions of the key words.

The four-step definition method can also help us to avoid the **strawman mistake**. In a discussion, it is very important to be as accurate as we can in representing another person's position or proposal. Just as we want them to take our views and ourselves seriously and not to distort or trivialize them, so we have an obligation to try as hard as we can to correctly and fairly formulate another person's position or proposal. When we fail to do this, when we distort or trivialize or otherwise misrepresent another person's view, we have made the Strawman Mistake, so named because it is usually easier to take apart (in the sense of criticize) a strawman than a real man. Sometimes this mistake is unintentional. People's views are often pretty complex, especially on important topics, and this means that it will be tricky to find a neat concise way to state them accurately. Just think of how complex your own views on, say, abortion, the desirability of a universal health care system, or of the justice of rising student tuitions. It would be hard for you to find a neat concise way to put your views. You want others to take as much care in formulating your views, as you would take in formulating them. You want them to use something like the four-step definitional method in stating your views, before they raise objections or criticism or state their own position. But sometimes, in debates and discussions, people deliberately distort their opponent's views, in order to gain a dialectical upper hand. This is always rude and boorish. But, even this makes it terrible at least from a critical thinking point of

view, distorting another person's view is always counter-productive. For by distorting or trivializing it we lose the opportunity to see whether there is any truth in it, and truth is after all one of the key goals of critical thinking.

5.4 REASONING ABOUT CAUSAL CONDITIONS

Trying to figure out what caused some event or what produces a certain kind of phenomena involves reasoning about the necessary or sufficient conditions that brought about that event or phenomena. If we want to know what caused little Joan's fever, what makes unemployment rise and fall, why some restaurant chains succeed and others fail, or why some hard-boiled eggs have dark coloring around the yolks and others don't, we need to reason about necessary and sufficient causal conditions. A causal condition is simply a condition that is necessary or sufficient for the occurrence of some event or phenomena. We have many ways to formulate claims about causal conditions. Here are some examples:

If you take ibuprofen your fever will subside.

Salmonella causes illness.

Smoking causes lung cancer.

Water helps to make grass grow.

Teasing your sister will make her angry.

Inflation produces unemployment.

We do not always formulate claims about causal conditions using conditionals, though as the example about ibuprofen illustrates, we can. We sometimes use the "causal" formulations as in the second and third examples. But we can also make causal claims using causal verbs, like "to make" and "to produce." In each of these cases, something is being said to be either necessary or sufficient for the occurrence of some event or phenomena.

Every event is caused by a condition that is (or by a set of conditions that are) sufficient for that event's occurrence and that condition or set of conditions must include everything that is necessary for the occurrence of that event. Some particular set of economic or financial conditions caused the recent rise in unemployment and this particular set is sufficient for the rise and it includes everything that is needed for unemployment to rise. Something happened to make the cake end up so bready, and those conditions are sufficient for the breadiness and

include everything that is required or needed for cake to end up breadly. Every set of sufficient causal conditions for some effect also includes every condition that is necessary for that effect. And just as with other necessary and sufficient conditions, a set of sufficient causal conditions need not be necessary (though it must contain every necessary causal condition), and a necessary causal condition need not be sufficient.

BOX 5.4 SUMMARY: CAUSAL CONDITIONS

Whenever some event occurs, all the conditions that are necessary for it also occurred, and some set of sufficient conditions also occurred.

5.4.1 The Meaning of Causal Claims

Unfortunately, we use the word “cause” in English to talk about both necessary and sufficient causal conditions. It is natural to understand the following as stating or identifying a sufficient condition:

Drowning causes death.

Drowning is not the only way to die. There are lots of ways to make something die, but drowning will do. Drowning is a sufficient causal condition for death, but (unfortunately, since drowning is easily avoided) not a necessary one. But here is a sentence that is naturally taken to identify a causal condition that is necessary but not sufficient:

Watering your lawn will make it grow.

Water is not all that a lawn needs to grow. Sunshine and the proper nutrients are also needed. But without water, the lawn will die. The presence of water is thus a necessary, but not a sufficient, causal condition for lawn growth. These examples show that we cannot always tell, just by the occurrence of the word “cause,” whether it is a sufficient or a necessary causal condition that is being identified. Still, it is usually pretty easy to tell which kind of cause a person has in mind when they make a causal claim. But we can avoid this ambiguity altogether if instead of using the word “cause” we use conditionals:

If a thing drowns, then it will die.

A lawn will grow only if it has water.

Once we formulate these claims about causal conditions using conditionals, then there is no ambiguity. The first one only states that drowning is sufficient for death, not that it is necessary, and the second one only states that water is necessary for lawn growth, not that it is sufficient.

BOX 5.4.1 THINKING CRITICALLY IN HISTORY

Sometimes, we can discover the causes of an event by thinking about what would have happened if things had been different. This reasoning involves the use of contrary-to-fact (or “counterfactual”) conditionals, like:

If Hitler had not invaded Poland, World War II would not have happened.

If humans had not emigrated from Africa thousands of years ago, there would not be humans now in North America.

If inflation had not spiked, the recession would have been less severe.

I'd be living in France today, if only I'd won that lottery.

In using a counterfactual conditional, one asserts (or at least, assumes) that the event specified in the antecedent did not occur, and one asserts that if it had occurred, the consequent would have been true too.

As the examples suggest, reasoning about counterfactuals is extremely common in the study of history. Thinking about how things might have turned out had the past been different in certain ways is one method historians use to discover the actual causes of actual events. (It is also very common in science.) Of course, knowing whether a counterfactual conditional is true can be just as difficult as knowing whether an indicative conditional is true.

5.4.2 Reasoning with Causal Claims

Reasoning with causal claims is just reasoning about necessary or sufficient conditions of a certain kind, and the valid forms of reasoning we discussed can be used when we reason about causal relations. The following are thus valid:

Taking ibuprofen will lower your fever, and little Joan just took some ibuprofen, so her fever will soon subside.

If inflation rises, then so does unemployment. But unemployment is holding steady, so inflation must not be rising.

If you over-stir a cake batter the gluten will begin to transform, and if the gluten begins to transform, the cake will be bready, so if you don't want a bready cake, don't over-stir the batter.

To reduce teen pregnancy rates, we need to change people's motivations, and we can do that only by spending more on targeted education. So if we are serious about reducing the pregnancy rate among teens, we need to spend more on targeted education.

The forms of reasoning that are valid for reasoning about necessary and sufficient conditions are the same whether we reason about causal conditions or about conditions that derive from definitions and stipulations. But unlike necessary and sufficient conditions that derive from standards and definitions, claims about causal conditions state matters of fact that we need to discover. We cannot (usually) just look in a dictionary or consult a rulebook to decide whether some causal claim is true. Whether rising inflation does cause unemployment is something we need to study. The factors affecting teen pregnancy rates have to be discovered through study. Usually, the causes of an interesting or important phenomenon are complex and nonobvious. But there are some strategies we can employ for deciding when a causal claim is acceptable or true.

An argument that has a false causal conditional might commit the **Slippery Slope Fallacy**: it is a mistake to argue using a causal conditional that is false or unjustified. Here is an example:

We should not let the county go ahead with plans to close the mental hospital. For, if the hospital closes, then all of its in-patients will be let out on the streets. And if that happens, then the crime rate will rise sharply. This would be unacceptable. We have to keep the mental patients in the hospital.

The author of this argument makes claims about the effects of closing the county mental hospital. There is, he is suggesting, a cascade of negative effects that would follow, creating a kind of slippery slope, and the effects would be so negative that avoiding them justifies keeping the hospital running. Suppose that those causal claims are false or unjustified. Suppose that either those effects would not follow (because, perhaps, the patients would be transferred to another hospital, or the

impact on crime of their release would be negligible) or that the author has no good reason to believe that they would follow. In that case, the author would be committing the slippery slope fallacy, by reasoning with a false or unacceptable causal conditional.

BOX 5.4.2 MISTAKES TO AVOID: SLIPPERY SLOPE REASONING

It is a mistake to reason with a false or unacceptable causal claim. This kind of mistake sometimes occurs in assessing policies or proposed actions, when someone argues that adopting the policy will lead to bad (or good) consequences, and either the causal claim is not justified or else the claim that the consequences are bad (or good) is not justified.

This example involves arguing that very negative things would be caused by some event, but the same mistake could involve arguing without sufficient justification that very positive things would happen.

We should all stop driving our cars to work. If we did this, the pollution in our skies would drop by 80 percent, the rivers would become clean once again, and fewer children would be killed in car crashes.

In the absence of very good reason to think that these positive would be the effect of a ban on driving this author commits the fallacy of slippery slope.

5.4.3 When Are Causal Claims Acceptable?

Finding out what causes what is usually extremely difficult. Just think of how long scientists and doctors have been trying to identify the causes of cancer. We have made a huge amount of progress, and we have been able to identify the causes of certain kinds of cancers, but a lot of expensive and sophisticated research remains to be done. But while the research and investigations needed to decide whether a causal claim is true are often time consuming and difficult, the methods involved in figuring out what causes what are, at least in the abstract, pretty straightforward, especially if we keep in mind the difference between a necessary and a sufficient condition. What makes it rare to find conclusive proof of the causes of an event is not that we lack a clear understanding of how to go about finding the proof—at least in

the abstract. Rather, what makes it difficult is that we usually lack the evidence that our methods tell us that we need to find.

5.4.3.1 Discovering Necessary Causal Conditions To identify the causal conditions that are necessary for some event, all we need to do is to see what is present every time the event is present. The idea is that whatever is necessary for the effect must be present whenever the effect is present. So, to find the necessary cause, look for the common factor. John Stuart Mill called this the **Method of Agreement**. Here is an example:

After eating lunch at the same restaurant, five individuals became ill with hepatitis. What caused the illness? Inspectors from the Health Department learned that while the five individuals had eaten different meals, they had all eaten tomatoes in their salad. Furthermore, the investigators found that this was the only food that they had all eaten. The inspectors concluded that the tomatoes had probably caused the hepatitis infection.

The researchers found an element that was common to all the cases: the eating of the tomatoes. The fact that tomatoes were eaten in all the cases where the illness was present is evidence that eating the tomatoes was a necessary cause of the illness, not that it was a sufficient cause. At most it shows that eating the tomatoes was a necessary part of the whole story of the illness. It does not show that other factors were not also involved. Maybe the tomatoes worked together with various other things, some of which were present in some cases and others in other cases, and that the tomatoes alone would not have caused it. So, the most that this evidence could show is that the tomatoes were a **necessary** part of the story. It could not show that they are the whole story.

Moreover, the fact that all of those who got sick ate tomatoes is not conclusive proof that the tomatoes are a necessary causal condition. The fact that they all ate the tomatoes might just be a coincidence. There are at least three ways this could turn out to be the case:

Overlooked conditions: The method of agreement yields conclusive evidence of a necessary causal condition only if we have identified all the possible similarities among the effects. But it is difficult to know for sure that we have. Perhaps those who got sick also shared a drink of beer. Because there are always so many different factors present in any series of cases, there can rarely be a guarantee that every single common factor has been identified. More-

over, the vast majority of the common factors will have nothing to do with the effect. The more certain we are that we have found all the similarities among the effects, the more certain we can be in our conclusion.

Multiple causes: The method of agreement yields conclusive evidence of a necessary causal condition only if the same cause is responsible for each occurrence of the effect. Perhaps both the potatoes and the liver were infected, and not the tomatoes, and that some of those who were infected were the only ones who ate the potatoes and the rest were the only ones who ate the liver, and it was just a coincidence that they all ate the tomatoes. It is hard to know whether all of the cases had the same cause without knowing what each of their causes was. But that is precisely what we are trying to find out.

Overlooked effects: The method of agreement yields conclusive evidence of a necessary causal condition only if one has identified all the cases where the effect is present. Maybe someone who got sick did not report it, and was not included in the study, and perhaps that person did not have the tomatoes. This would show that eating the tomatoes was not a necessary cause at all.

The more certain one is that one has examined all the similarities among the case where the effect is present, that one has ruled out possible multiple causes, and that one has studied all cases of the effect, the more certain one can be in one's conclusion about the effect's necessary causal conditions.

5.4.3.2 Discovering Sufficient Causal Conditions To identify the causal conditions that are sufficient for an event, look for what is missing whenever the effect is missing, but sometimes present when the effect is present. The idea is that something sufficient for the event will be absent whenever the event is absent and whenever it is present the effect will be present too. John Stuart Mill called this the **Method of Difference**.

After conducting a study on the work force at a certain factory, engineers found that five of their 20 workers performed their task less efficiently than the others. A list was made of the various factors that were present in the case of the fifteen productive workers, but absent in the case of the five unproductive workers. It was discovered that among eight likely candidates, only one factor was missing for all five: they all complained about not getting any share of the profits. The engineers learned that

many of the other workers with high productivity did participate in the company's profit sharing plan. They concluded that having a share in the profits increases worker efficiency.

This evidence shows at most that profit sharing is one way to increase worker efficiency, not that it is the only way. Maybe there are other ways to achieve the same effect. What is more, and just as with the method we studied for identifying an effect's necessary causal conditions, this method yields conclusive evidence of a sufficient causal condition only when (i) no missing conditions were overlooked (maybe something they did not notice was also missing in all five cases); (ii) there is only one sufficient cause at issue (maybe two of the unproductive workers would respond to higher pay while the other three would respond to shorter hours); (iii) no cases were overlooked (maybe a sixth unproductive worker does not care about profit sharing at all.)

5.4.3.3 Discovering Necessary and Sufficient Causal Conditions

The following example illustrates how we can combine both these strategies to try to identify conditions that are individually necessary and jointly sufficient:

Eight inhabitants of a town contract a rare form of the plague. A doctor is flown to the town with a serum she thinks might be a cure. Only four of the inhabitants accept the cure, the other four insist on using home remedies. But all eight had been treated with home remedies before the doctor arrived. Eventually, the four who received the serum recovered, while the other four died. What had caused the recovery? The doctor noticed that among those who survived, no single home remedy was given to all; and that each home remedy had been given to at least one of those who did not survive. The doctor concluded that the serum caused the recovery.

Here the doctor's conclusion is that the serum is necessary for a cure (since it was present every time the cure was present) and sufficient too (since it was absent every time the cure was absent.) As before, though, this evidence is not conclusive, for all the reasons we have noted.

5.3.4.4 Concomitant Variation When two or more phenomena vary together—what we call **concomitant variation**—then this is some reason to think that there is a causal link involved. If the rise in interest rates is always accompanied by a rise in unemployment, then

this is evidence that a causal process is responsible for the correlation. If a lowering of a lake's water temperature is correlated with an increase in a frog population, and that population decreases when the temperatures rise, then this too is some reason to think there is a causal link involved. Likewise, if changes in one phenomenon can happen without any changes in another phenomenon, then this is some reason to think the two phenomena are causally independent or isolated from one another. If changes in the blood sugar level of white mice occur without any change in the rate of spread of their cancer cells, then this is some reason to think there is no causal link between the two. If the amount of mineral residue is the same even when there are changes in the amount of soap used in a restaurant dishwasher, then this is some reason to think that the soap is not causally relevant to the mineral deposits. But the mere existence of concomitant variation among phenomena does not tell us much about the causal link. It might be that changes in one of the phenomena are causing the changes in the other. Or it might be that changes in the observed phenomena are being produced by changes in some underlying phenomena.

5.3.4.5 Experimenting and Simulating **Experimentation** can help us in our reasoning about necessary and sufficient causal conditions. As we saw, the methods of agreement and difference usually do not provide conclusive evidence because there are usually too many commonalities to know for sure which ones are causally relevant and which are merely coincidental. Only in very special and rare cases can we get conclusive evidence of an event's cause. This limitation is especially troublesome when, as in the examples we discussed, only a very small number of very similar cases are studied. For the smaller the number of cases studied, and the less diversity there is among them, the larger the number of commonalities, both of those things present and of those things absent. What we need is a way to generate a greater number and variety of cases to study. This is where experimentation comes in handy. When we set up an experiment, we can create as many cases to study as we can afford, with as much variety among them as can we imagine.

Suppose you want to figure out what causes some phenomena, P. The first thing you need to do is to make a list of possible causes. You then have to design an experiment that will show, for each possible cause, whether it is necessary or sufficient. It is very important that only one possible cause be tested at a time. That is, you need to vary only that one condition; you must **control** for variations in the other

conditions. If your initial hypotheses were right, and if your experimental controls were effective, then you will gain additional evidence of a cause. If not, then you have to reconsider either the hypotheses or the experimental design.

Suppose the researchers in the restaurant case wanted to set up an experiment to test their hypothesis that eating the tomatoes caused the illness. Suppose, in particular, that they want to rule out the possibility of multiple causes. (Notice the way that this would involve reasoning about alternatives.) They want to test their hypothesis against the possibility that some of the people got sick from oysters and the others from shrimp. The researchers might take samples of all of the other food and drink that those who got sick had eaten and feed them to laboratory mice and see whether any of them get sick. If none do, then that is some reason to think that eating the tomatoes was indeed a necessary part of the cause of the original illness. To see whether eating the tomatoes was also a sufficient cause, they might try feeding just the tomatoes to a group of mice, and see whether any of them get sick. If they do, then that is some reason to think that the tomatoes alone were the cause. If the mice fed the tomatoes do not get sick, the researchers might try combining tomatoes with other common elements to see whether some combination of the foods produced the illness. By creating more cases to study, and by making them as varied as possible, the researchers can gain additional evidence. The very same methods of looking for commonalities either present or absent are then used to draw conclusions from the newly gathered evidence. Experimentation can help us to discover the causes of things by providing new cases to study with the old methods.

BOX 5.4.3A DECIDING WHAT TO DO: COMPARING CONSEQUENCES

A crucial step in deciding what to do is identifying the likely consequences of various courses of action. This involves reasoning about causal conditions, about what would happen if one adopted each of the candidates under consideration. The negative effects of a course of action are that action's potential **costs**, while the positive effects are its potential **benefits**. Knowing the likely costs and benefits of a proposal usually requires knowing a lot about the case at hand. The methods we have studied in the text can help you figure this out.

In addition to considering the direct costs and benefits of a proposal, it is also important to consider the indirect costs, including the costs associated with lost opportunities. Adopting one course of action over another will impact what opportunities are available at a later date. (Buying a stereo now means not buying a government bonds tomorrow.) In effect, in adopting a course of action, one would be giving up the benefits of those opportunities. They would be “lost opportunities.” Any benefit from lost opportunities should be counted as an indirect cost—an “opportunity cost”—in assessing the overall effect of a course of action.

Researchers can sometimes use **simulations** when it is not possible or practical to produce actual cases for study. Researchers studying extreme weather events, such as hurricanes and tsunamis, cannot just to go into their lab and produce hurricanes and tsunamis to study. But they can create computer models that simulate weather conditions and run them to see how different factors affect the resulting weather patterns. Economists create models of financial markets and study how changes in one element in the market can change other elements. The military runs virtual “war games” to see how different tactics affect the battlefield. This kind of experimentation relies on the use of models, and the value of the evidence they provide of causal conditions depends on whether these models are accurate representations of the phenomena they are trying to study. In Chapter 6, we will look in more detail at the use of models in reasoning.

A common mistake in reasoning about causal conditions is to conclude from the fact that one event follows another that the first is causally connected to the second. This mistake is called the **Post Hoc Fallacy**. While it is true that the causes always precede the effect, this is just a necessary condition for a causal link not a sufficient one. Indeed, as our examples have illustrated, much more evidence is needed than merely the temporal order of two events before one can draw any conclusions about the causal relations between them. One sometimes hears politicians championing the virtues of their tax cut plans by noting that unemployment fell, government revenues rose, and the general level of happiness soared to an all-time level, right after their favored tax cuts went into effect. Such claims should always be taken with a grain of salt. The historical order of those events is quite simply insufficient evidence on which to base such a claim. (Moreover, those politicians are not trustworthy witnesses: they are probably not properly trained or informed in the economics of tax cuts and they are probably biased.)

Of course, as always, they might be right: a true belief might be based on inadequate or unacceptable evidence. But if this is the sole basis for their belief, then that belief is unjustified.

BOX 5.4.3B MISTAKES TO AVOID: POST HOC REASONING

A common mistake in reasoning about causal conditions is to conclude from the fact that one event follows another that the first is causally connected to the second. This mistake is called the **Post Hoc Fallacy**. While it is true that the causes always precede the effect, this is just a necessary condition for a causal link not a sufficient one.

EXERCISE 5.4

1. Short-answer questions:

- a. What is the difference between a necessary and a sufficient condition?
- b. Give five examples of a necessary causal condition for some effect that are not sufficient conditions for it, and five examples of a sufficient causal condition for some effect that are not necessary for it.
- c. Could something be a necessary element of a sufficient causal condition for some effect without being a necessary causal condition for that effect? If so, give five examples.
- d. Why is experimentation not a new method?

2. Which method of reasoning about causal conditions is being used in the following cases? Using the concepts we have been discussing, describe five factors that make the evidence less than conclusive. Be as specific as you can.

- a. Mary noticed that all the cookies that had been kept in the plastic bags were dry and those kept in the freezer or in a cookie jar had stayed fresh. She concluded that keeping cookies in a plastic bag makes them stale.
- b. The city noticed that many of the parking meters were broken, and found that all of the broken ones had been tampered with by having bubble gum inserted into their coin slots. The city concluded that the gum was causing the problem.
- c. Doctors at the county hospital noticed that many of their patients returned after a few days complaining about a skin rash. They

did a quick survey of those complaining of the rash, and found that all of them had used the hand soap dispenser outside the entrance doors. The doctors concluded that something in the hand soap was causing the reaction.

- d. The local radio station noticed that during the hours when they played only classical music their listenership went down, and that it went up again when they switched to bluegrass music. They decided to become an all bluegrass station in order to maximize listeners.
 - e. Voting just does not make a difference to what the government does. In every democracy, it is still big business interests that decide government policy. And in nondemocratic countries, it is still big business interests that decide what the government does. So voting makes no difference!
 - f. The drug company Pharmastock did a study of its new anti-obesity drug. It found that those who took the drug regularly lost more every week than those who took a placebo (i.e., a tiny sugar pill). It kept a careful watch to make sure that nothing else was different between the two groups. Pharmastock reported to its shareholders that its new drug was a huge success.
 - g. Susan planted six rows of corn in her garden. Along two of them she also planted cone flowers. Along two others she planted miniature rose bushes, while along the last two she planted nothing at all. The plants all grew very well, but during the harvest she noticed that the ears of corn on the last two rows had all been eaten by bugs, but that the rest were fine. She decided that planting the flowers prevented bug infestations.
3. What conclusions about the necessary or sufficient causes could be drawn from the following evidence? Using the concepts we have been discussing, explain why your answer is right, but also why those conclusions are not conclusive.
- a. Doctors at Harvard who studied 2,000 nurses for 25 years found that those who had high cholesterol came from different ethnic backgrounds and different economic classes.
 - b. Studies on tadpoles found the following. If the tadpoles were spawned in pond water that averaged more than 75 degrees, the tadpole population had a larger percentage of mutations and if the water was on average below 65 degrees there were no mutations at all. However, if the temperature varied from a high of 75 to a low of 65 then there were also no mutations.

- c. Four groups of pregnant mothers were studied. The first group was given a shake that was high in protein but low in iron every breakfast. The second group was given a shake low in protein but high in iron. The third group was given a shake high in both protein and iron. The final group was given a shake that contained no iron or protein. After three weeks, the women in the third group had more stable blood sugar throughout the day than the women in other three groups.
 - d. Partners at a local marketing firm did a study at 20 local shoe stores. The firm found the following: stores that played music with a slow beat had the worst sales of all the stores, even worse than those that played no music at all. Those that played music with a fast beat had the best sales, and among those, the ones that played it loudest had the best sales. The researchers found no other commonalities among the different groups of stores.
4. Describe an experiment to test the following causal claims. Using the concepts we have been discussing, explain how the test is designed to answer the question.
- a. Adding speed bumps to residential city streets reduces the number of accidents.
 - b. Adding mustard to a salad dressing makes the oil and vinegar stay blended longer.
 - c. The more a runner stretches before her long runs the fewer injuries she gets.
 - d. High temperatures make ladybug populations decline.
 - e. Kneading bread dough for more than ten minutes reduces the number of air holes in the finished loaf.

5.5 MISTAKES TO AVOID

False Alternative It is a mistake to reason about alternatives with a false disjunction. Not only does this make the disjunction an unacceptable premise, but also it guarantees that the conclusion of the reasoning—even if that reasoning is valid—will be false.

Lucky Disjunction It is mistake to reason about alternatives with a disjunction that one knows (or should know) is not exhaustive. Such a disjunction might still be true—but one would not know it to be true. One’s reasoning would reach a true conclusion only by luck, and it is wrong to rely on luck over careful thought.

Affirming a Disjunct It is a mistake to conclude from the fact that one disjunct of a disjunction is true that the remaining disjuncts are false, unless one knows that the disjunction is exclusive.

Denying the Antecedent It is a mistake to conclude that the consequent of a conditional is false just because the antecedent is false. An argument of this form is not valid: If P, then Q; it is not the case that P; so, it is not the case that Q. This mistake involves confusing a sufficient condition for a necessary one.

Affirming the Consequent It is a mistake to conclude that the antecedent of a conditional is true just because the consequent is true. An argument of this form is invalid: If P, then Q; Q is true; So, P is true too. This mistake involves mistaking a necessary condition for a sufficient one.

Red Herring Fallacy It is wrong to introduce something irrelevant in a discussion in order to divert attention away from an author's views or reasons. We saw that this can occur when, in the assessment of an author's reasoning about alternatives, a possibility that is unreasonable or highly unlikely is raised to show that an author is relying on a disjunction that is not exhaustive. Enough is enough; it is not everything.

Appeal to Ignorance It is wrong to believe or decide that something is true (or false) just because it has not been proven otherwise. This can occur when reasoning about alternatives if one believes or judges that a disjunction is true or exhaustive just because one does not know of or cannot think of any other possibility. It also happens when one concludes that a claim is true on the grounds that one has not found a counterexample to it.

Strawman Mistake It is wrong to distort or misrepresent another person's views or reasons. Not only does this hurt their feelings (and when done deliberately is exceedingly boorish), but it undermines the goals of critical thinking, which is to seek the truth in another person's views, even if those views are not our own.

Equivocation It is a mistake if a word or phrase in an argument has to mean one thing for the argument to be valid and something else for the premises to be true.

False Conditional (Slippery Slope) Fallacy It is a mistake to reason with a false or unjustified causal conditional. It is, of course, always better to ensure that the premises of one's arguments are acceptable.

Post Hoc Fallacy It is a mistake to conclude that one event caused another solely on the grounds that the first preceded the other. While it is true that a cause always precedes its effect, this is not a sufficient condition for causation.

5.6 PRACTICAL STRATEGIES

Counter-Example Strategy A counter-example to a claim is an example that shows that that claim is false. The example has to be one that everyone in the discussion can accept; otherwise it will simply beg the question. Finding a counter-example proves that the claim is false; failing to find one, though, does not show that the claim is true.

Method of Agreement To find a necessary causal condition for some effect, look for something that is present whenever the effect is present. The method relies on the idea that everything necessary for the occurrence of some effect will be present whenever the effect is present.

Method of Difference To find a sufficient causal condition for some effect, look for whatever is absent when the effect is absent and that whenever present is followed by the effect. The idea is that whenever a sufficient condition is present the effect will be present too.

Method of Concomitant Variation If two phenomena vary together, then this is some evidence of a causal relation between them. But by itself, this does not reveal much about the link. One of the phenomena might cause the other or they might have a shared underlying cause.

Experimentation and Simulation Knowing for sure whether some condition is a necessary or sufficient cause of some effect requires ruling out a lot of alternative possible causes. Doing this requires accumulating more evidence. New evidence can be generated by producing real (through experimentation) or fictional (through simulation) cases of the phenomena in question. In both cases, it is best to try to vary one possible cause at a time. This is called "controlling".

6

REASONING BY ANALOGY

Critical thinking is reasonable and reflective thinking aimed at deciding what to believe and what to do. Part of what makes it reasonable is that it insists that we have sufficient reasons for our beliefs and decisions. In Chapter 3 we saw that a **valid** argument provides an ideal model of reasons that are sufficient: an argument is valid when it is not possible for the premises to be true and yet for the conclusion to be false. When an argument is valid, the truth of the premises guarantees that of the conclusion. In that case, the reasons are logically sufficient for the conclusion. But we also saw that relying on the ideal of a valid argument in order to evaluate reasons is sometimes difficult, since deciding whether an argument is valid requires us to assume for the sake of an argument things we might not be quite ready to believe and then to ask whether something else would have to be true given that set of assumptions. Thankfully, a good deal of our reasoning can be done using forms of reasoning that are guaranteed in advance to be valid. When we employ these valid forms of reasoning our only remaining question in evaluating whether to accept the conclusion is whether the premises are true.

In Chapter 5, we studied reasoning about alternatives and about necessary and sufficient conditions. We saw that any argument of the

form of **modus ponens**, **modus tollens**, or **denying a disjunct** is a valid argument. In this chapter, we will study **reasoning by analogy**. Reasoning by analogy involves drawing conclusions about one thing or kind of thing by comparing it to something else that is like it in relevant ways. As we will see, this kind of reasoning is very common and very powerful. It is basic to a good deal of our ordinary, common sense reasoning about what to believe and what to do. But it is also basic to a good deal of the reasoning in the natural and social sciences. We will first study reasoning by analogy in the abstract and then look at some particular applications of it. But let's start with some examples, to get a sense of just how varied and common this kind of reasoning is.

6.1 REASONING BY ANALOGY

6.1.1 Examples

Here are some examples of reasoning by analogy:

1. John is just like his brother Peter, and Peter is a really generous guy, so I think that John must be generous too.
2. The war in Iraq is just like the war in Vietnam. We lost in Vietnam because we left too early. So, we should not leave Iraq until the war is won.
3. When you do chemistry, you have to be really careful and precise in your measurements and timing. And baking bread is just like doing chemistry. So it is important to be careful and precise when baking bread too.
4. Eighty percent of those we surveyed believe that we should build a new bridge across the river instead of a tunnel. Clearly, the majority of the city prefers the bridge idea.
5. According to the computer models, the storm will continue to track north for another two hours, and then move east into the direction of the city. So, the city is probably going to get a big rainstorm in about two hours.
6. The map says the buried treasure should be right here. Start digging!

The first three arguments pretty clearly involve reasoning by analogy. Each has a premise that compares two things. In the first argument, John is compared to his brother; in the second argument, the war in Iraq is compared to the war in Vietnam; in the third argument, baking

bread is compared to doing chemistry. And in all three arguments, a conclusion about one of those two things is supposed to follow from this comparison or analogy together with an additional claim about the other thing. The reason to think that John is generous is that he is (according to the arguer, anyway) just like his brother Peter and Peter is generous. In the second argument, the (supposed) fact that the war in Iraq is just like the war in Vietnam, together with the (supposed) fact that the war in Vietnam was lost because U.S. troops left too early is supposed to show or prove that the U.S. will lose the war in Iraq if it leaves too early. In the third case, the reason to be careful in measuring when baking bread is that (according to the argument, anyway) baking bread is just like doing chemistry and doing chemistry requires careful measurement.

Before we consider the remaining three examples, see the Box titled, “Summary: Reasoning by Analogy” for three clear cases to identify the abstract form of reasoning by analogy.


BOX 6.1 SUMMARY: REASONING BY ANALOGY

The form of reasoning by analogy is this:

(1) X has property P.

(2) Y is just like X in all the respects relevant to P.

(3) So, Y is P too.



Some terminology will help. Let’s call the premise that states the analogy or comparison the **analogical premise**. Let’s call the two things that are being compared the **analogues**. Finally, let’s call the property claimed to be true of the conclusion’s subject the **relevant property**.

6.1.2 Is Reasoning by Analogy Valid?

We know that there are two very different questions to ask when evaluating an argument: are its premises true, and do they support the conclusion? We also know that these questions are independent of one another. Whether an argument’s premises provide logical support for the conclusion has nothing to do with whether they are in fact true. Rather, it is a matter of whether the conclusion would have to be true *if the premises themselves were* true. But this is a difficult distinction to keep in mind, and it is easy to mistake uncertainty about the truth of

an argument's premises for uncertainty about the argument's validity. The more we can do to keep these different kinds of uncertainty distinguished, the better we will be as critical thinkers. None of this is new; we have been making this point over and over in this book. But it is of special relevance in this chapter, since it is often difficult—even for the best critical thinkers—to know whether the problem with an argument by analogy is that the premises are unacceptable or whether they are insufficient (or both!). In this section we will make a start at learning how to draw this distinction.

Let's start by considering the first example from above. Let's recast it in our abstract form, as follows:

1. John is just like Peter in all relevant respects.
2. Peter is really generous.
3. So, John is really generous.

And let's ask whether it is valid, using the test for validity we discussed in Chapter 3. Suppose (just for the sake of the argument) that the two premises were true. That is, suppose that Peter is a really generous person. And suppose that John really is just like Peter in all of the relevant respects. Could it nonetheless still be that John is not generous after all? Could it be that John is a tightwad, always refusing to leave a tip in a restaurant, say, or never offering to buy anyone a coffee or a cookie? How could it be? If John really is such a tightwad, then either he is not really just like Peter *in all relevant respects*, or else Peter himself is not really generous. In other words, if John is not generous, then either Peter is not generous, or else there is some *relevant* difference between Peter and John that explains why one is generous and the other is not. At least one of those premises would surely have to be false if the conclusion was false. We might not know which was false, of course. But if we knew for a fact that John was a stingy guy, then we would surely know that either he is not like his brother, or else his brother is stingy too. All of this shows, I think, that if the argument's premises were true, then its conclusion would have to be true too. And this just means that the argument is valid.

6.1.3 Relevant Similarity

It might have already occurred to you that in reaching the conclusion that the argument is valid, we relied pretty heavily on one central part of the analogical claim. I even put it in italics: *in all relevant respects*. The analogical premise says that John is just like Peter *in all relevant*

respects. What does this mean? What does it mean for them to be alike in all relevant respects? Relevant to what? In this case, to generosity, of course. Which respects are relevant in an argument by analogy depends on the **relevant property**. It might be that John is just as tall as Peter, or just as good at drawing. But these similarities are not relevant to generosity. Whether one is generous has nothing much to do with one's height or one's artistic abilities. But the fact that John is just as empathetic as Peter, and just as kind as Peter and has the same keen sense of duty and justice as Peter, are similarities that are relevant to generosity. Whether a person is generous does depend on how well she can see things from another person's point of view (and so on how empathetic she is) and how keenly aware she is of her moral and social obligations. No doubt there are other elements that are also relevant to generosity. In claiming that John is just like Peter *in all relevant respects*, the analogical premise means that the two are alike in all the ways that are relevant to whether a person is generous or not. If Peter is very empathetic, and if being empathetic is relevant to generosity, then, that premise says, John is just as empathetic as Peter.

Now that we know the intended meaning of the phrase "in all relevant respects," the validity of reasoning by analogy should seem not very surprising. If such and such really is true of X, and if X and Y really are just alike in all the respects that are relevant to whether something is such and such, then such and such would surely have to be true of Y too. There is no other way. If such and such is not true of Y, then either it is not really true of X after all, or else there is some relevant difference between X and Y that explains why it is true of one but not the other. This confirms our earlier conclusion that reasoning by analogy, when it takes the abstract form we identified, is indeed a valid form of argument.

Notice that in this form of argument the **relevant property** mentioned in the conclusion is also mentioned in the premise about the other analogue. This is crucial to this form's validity. In the following argument, this is not the case:

Peter is very generous, and John is just like him in all the relevant respects, so John is a nice guy too.

The relevant property in this argument is being a nice guy. Being generous is not exactly the same as being a nice guy. Probably, being nice requires being generous, but perhaps someone who is not generally very nice might still be generous—their one redeeming feature.

An argument by analogy could be valid even if the two claims are not exactly the same, *so long as the one in the premise includes or implies the one in the conclusion*. Here is an example:

The new sofa is just like that old one, and the old one fell apart within six months, so this new one will probably not last long either.

The premise claims that the old sofa did not last six months, whereas the conclusion claims that the new sofa will not last long. These are not really the same claims. But the one in the conclusion surely follows from the one in the premise: if a sofa does not last six months, then it does not last long. In effect, this conditional is a missing premise in the argument: adding it in helps us to see that the argument is valid.

6.1.4 When Is an Analogical Claim True or Acceptable?

As we have now seen time and again, the fact that an argument is valid does not mean that its conclusion is true. Nor does it mean that we yet have good reason to believe its conclusion. Whether an argument constitutes good reason to believe its conclusion also depends on whether the argument's premises are true or acceptable. In the case of an argument by analogy, this comes to the question of whether the analogical premise is true or acceptable. In this section, we will study how to tell whether an analogical premise is true.

Deciding whether an analogical premise is true is difficult for two related reasons. An analogical claim asserts that the analogues are alike in all relevant respects. The first difficulty is that any two things in the universe are alike in lots and lots of ways. (There may even be an infinite number of similarities between any two things!) Likewise, there are lots of ways in which any two things are different, perhaps even an infinite number of them. So, deciding whether some analogical claim is true is not simply a matter of deciding whether the analogues are alike: we already know that they are. Crucially, of course, the analogical claim is also asserting that the analogues are alike in all *relevant* respects, and this is what introduces the second and greatest difficulty in deciding whether the analogical claim is true. For it is not always easy to know what the relevant respects are, let alone whether the analogues are alike in them. Take our example of John and his brother. In deciding whether it is true that they are alike in all respects relevant to generosity, we face two questions:

1. What respects are relevant to generosity?
2. Are John and Peter alike in such respects?

Or consider our second example, the one in which the war in Iraq is compared to the war in Vietnam:

1. What respects are relevant to losing or winning a war?
2. Are the wars in Iraq and Vietnam relevant in those respects?

Let's call the first kind of question the **relevancy question**—since it asks which respects are relevant, and let's call the second kind of question the **comparison question**—since it asks whether the analogues are alike in those respects.

BOX 6.1.4A SUMMARY: ACCEPTABLE ANALOGICAL CLAIMS

To assess whether an analogical claim is true, two questions need to be asked:

Relevancy question: which respects are the relevant ones?

Comparison question: Are the analogues alike in the relevant respects?

Relevancy questions are very difficult to answer. In some ways, answering a relevancy question it is like trying to decide whether a claim about necessary and sufficient conditions is true. As we saw in Chapter 5, deciding whether a conditional is true can require knowing a lot of factual information about the subject matter of the conditional. We cannot know, just by thinking about it, whether a plant will die, if it is deprived of nitrogen. We need to do studies and experiments to find out whether nitrogen is necessary for plants to live. The same is true in the case of the relevancy question and so in reasoning by analogy. Knowing what factors are relevant to winning a war or to being generous may require knowing a lot about wars and generosity, and this is not knowledge one can acquire just by thinking about it. We need to do some studies about wars and about human behavior in order to know what factors are relevant to winning a war or to being a generous person. This makes it difficult to evaluate the truth of analogical claims.

Comparison questions are also difficult to answer. Even if we did know what respects were relevant, it might still be difficult to know whether the analogues are alike in those respects. Suppose we learn that being empathetic is relevant to being generous. How can we tell whether John is empathetic? And how can we tell whether he is *just* as

empathetic as his brother Peter? How can we measure empathy? Suppose that lacking public support at home is relevant to whether a war is winnable. What is the best way to figure out whether the public really does support the war in Iraq? And how could we figure out whether the public support for that war is just like the public support for the war in Vietnam? This question can be just as difficult as the relevancy question.

An argument by analogy that has a false analogical premise commits the mistake we call **false analogy**. As we have been seeing, it can be difficult to know for sure whether an analogical claim is true. It is the responsibility of the person advancing the argument to show that the analogical claim is true or acceptable. But if we are assessing an argument by analogy, we should not accuse it of committing the mistake of a false analogy unless we can back up that accusation. More specifically, we should not accuse an argument of false analogy unless we can point out some relevant difference between the analogues.

BOX 6.1.4B MISTAKES TO AVOID: FALSE ANALOGY

It is a mistake to reason by analogy using a false or unacceptable analogical claim. Recall that any two things are alike in a huge number of respects. An analogical claim is true only if the analogues are **exactly alike** in all of the respects that are **relevant** to the relevant property. Knowing what those respects are can be difficult, and can sometimes require a lot of investigation. Knowing whether the analogues really are alike in those respects can also be difficult.

If you are **constructing** an argument by analogy to support some claim, then you have an obligation to make sure that the analogical claim is acceptable.

If you are **evaluating** an argument by analogy, and you believe that the analogical claim is not acceptable, then you have an obligation to identify some relevant respect in which the analogues are not alike.

We have been considering when an analogical claim is true or acceptable. Knowing whether one is true can be difficult, we have seen, both because it can be difficult to know which respects are the relevant ones and because it can be difficult to know whether the analogues really

are alike in those respects. Often, this uncertainty is expressed by adding a word like “probably” into the conclusion, as in the following example:

The model of the airplane was able to withstand strong cross forces, and the model is quite accurate, so the plane will probably be able to withstand those forces too.

It is always a good thing to be honest about how certain we are about the acceptability of an argument’s premises. Using words like “probably” to make clear our level of certainty is a good thing. But none of this uncertainty about the acceptability of our reasons would show that the argument is not *valid*. Uncertainty about the truth of an argument’s premises is not uncertainty about whether the truth of the premises would be sufficient for the truth of the conclusion. Since it is relatively easy to construct a valid argument, we should never be uncertain about the validity of our arguments. We can be certain that an argument is valid, even if we are not certain whether its premises are true. This is such an important point that it is almost impossible to repeat it too often.

BOX 6.1.4C CRITICAL THINKING AND THE LAW

Reasoning by analogy plays a crucial role in legal theory in at least two places.

- *Similar cases should be treated similarly.* Crimes that are similar in relevant ways ought, all things considered, to be punished in the same way, and ones that are dissimilar in relevant ways ought, all things considered, to incur different punishments. In the absence of a relevant difference, it would be unfair to sentence one tax evader to jail but another to probation. Likewise, it would be wrong to punish a jaywalker as severely as a mass murderer. Of course, the difficulty is in deciding which respects are the relevant ones.
- *Filling legal gaps.* No matter how carefully laws are written, there are bound to be unforeseen cases. Judges (and prosecutors and legislators) have to decide how to treat unforeseen cases. Often, the similarities between the unforeseen and the intended cases play an important role. Again, the difficulty is in deciding which similarities are the relevant ones. Given that

the U.S. Constitution says nothing about the right to own submachine guns, is the ownership of such weapons more like the ownership of tanks or more like the ownership of muskets?

EXERCISE 6.1

1. Short-answer questions:

- a. What are analogues? Use the four-step definition method from Chapter 2 to develop your answer.
- b. What is the form of an argument by analogy?
- c. Using the concepts of a necessary and sufficient condition, explain why an argument by analogy that fits the form is valid.
- d. Under what conditions is it logically acceptable for the property mentioned in the conclusion of an argument by analogy not to be the same as the property mentioned in the premises?
- e. When is an analogical claim true or acceptable? Use the four-step definitional method from Chapter 2 to develop your answer.
- f. What is the mistake of a false analogy and why is it a mistake? Use an example.
- g. Are the premises in an argument by analogy dependent? Use an example.

2. We have seen that analogues might be alike in respects that are relevant to one feature, but not alike in respects relevant to another feature. For each pair, find one feature with respect to which they are alike in all relevant respects and one feature with respect to which they are unlike in some relevant respect.

- a. A hockey game; a game of chess
- b. A garden; an economy
- c. A human; a mouse
- d. A car; the solar system
- e. A watch; the universe
- f. A toy train; a real train
- g. An electron; a planet
- h. Beef; tofu
- i. Loneliness; love
- j. Juliette; the sun

3. Rewrite the following arguments to make their form clear. Insert a missing analogical claim if one is needed.
- a. Love is like a garden and if you do not tend a garden every day the flowers will die. Love needs tending too.
 - b. I won't like that movie. It is just like that other one we saw last week, and I hated it.
 - c. Cinnamon is just like mace, and the mace tasted good in the cake, so the cake will taste good too with cinnamon.
 - d. The stock market is behaving just like it did two years ago, and back then we had a very deep recession, so another recession is on the way.
 - e. Crime rates are about to drop. That's what happened in Europe when the economy boomed, and our economy is booming too.
 - f. The samples taken from the patient's wound were infected, so the wound is now badly infected.
 - g. The sun has risen every day for nearly 6 billion years. Tomorrow will be no different. So, the sun will rise tomorrow too.
 - h. We repeated the tests 300 times and got the same results, so alcohol does kill germs on skin.
 - i. Demolishing vacant buildings is like pulling out weeds. It increases the value of the surrounding buildings.
4. Look in the letters to the editor in a newspaper or magazine. Find three letters in which the author argues by analogy. Identify the analogues, and assess whether the analogical claim is acceptable.

6.2 REASONING USING REPRESENTATIONAL ANALOGY

Let's return to the arguments we started with. The first three are obviously arguments by analogy, since each of them has an analogical claim as a premise. But, what about the other three examples?

1. Eighty percent of those we surveyed believe that we should build a new bridge across the river instead of digging a tunnel. Clearly, the majority of the city prefers the bridge idea.
2. According to the computer models, the storm will continue to track north for another two hours, and then move east into the direction of the city. The city is probably going to get a big rain-storm in about two hours.

3. The map says the buried treasure should be right here. Start digging!

These arguments do not have an analogical claim as an explicit premise. The first one has a premise about a survey of public opinion; the second one has a premise about a computer model of a storm; the third one involves a claim about a map. Still, in each case, a comparison of one thing with another—the surveyed group to the general population, the computer model of the storm to the real one, the map to the island—is an essential premise in the reasoning. This is so even though the comparison is not explicitly stated. The presence of this (albeit implicit) comparison is what makes them all cases of reasoning by analogy. They are just like the case of John and Peter, only the comparison is not explicit.

Still, there is an important difference between these three cases and the ones with which we began. In the first three cases, the analogues are things of the same general kind: John and Peter are both humans; the wars in Iraq and Vietnam are both wars; and baking bread and doing chemistry are both activities that require measurement. But in the case of these last two arguments, the analogues are things of very different kinds: a computer model of a storm is not a storm, and a map is not an island. It might sound odd to suppose that a map could be in relevant respects just like a desert island, or that a computer simulation could be in relevant respects just like a real hurricane. But in these cases, the comparison rests on the claim that one of the analogues is an **accurate representation** of the other. The map is assumed to be a pictorial representation (i.e., a kind of picture) of the island and the computer simulation is assumed to be a representational model of the real storm. The reasoning relies on this assumed representational relation between the analogues. If the map is an accurate representation of the island, then one can rely on it when deciding where to dig for the treasure. If the computer model is an accurate representation of the real storm, then one can rely on it when predicting how the storm will move. The very same is true of the argument involving the survey: the group of people surveyed is assumed to be **representative** of the entire population. If it is, then one can rely on what we know about it to draw conclusions about the entire group. Reasoning that draws a conclusion about something on the basis of a representation of it still counts as reasoning by analogy, and what we have said so far about reasoning by analogy applies to it as well. Let's look at each kind in turn.


6.2.1 Reasoning with Samples

Reasoning with samples involves drawing a conclusion about something on the basis of a claim about a sample of it. This kind of reasoning is extremely common. Here are just a few examples. To decide whether her soup needs more salt, a chef may taste a spoonful of it. The spoonful is a sample of the soup. To decide whether her patient has strep throat, the doctor might collect a sample of the infection on a long cotton swab. A diabetic tests a drop of his own blood to see whether his blood sugar is high. To find out what Americans think about the risks of global warming, pollsters might ask a randomly chosen sample of 2,500 Americans. A movie's producers invite a sample audience to watch a preview before deciding whether to make additional changes to it. All of these cases involve reasoning using samples. In this section, we will study when reasoning using samples is logically good.

BOX 6.2.1 SUMMARY: REASONING WITH SAMPLES

The form of reasoning with samples is this:

- (1) Such and such is true of X.
- (2) X is representative sample of Y.
- (3) So, such and such is true of Y, too.



The premise stating that X is a representative sample of Y is the **analogical premise**. Let's call X the **sample group** and Y the **target group**. Just as with reasoning by analogy, an argument of this form is valid: if its premises are true, then so is its conclusion. If such and such is not true of Y, but it is true of X, then it must be that X is not really a representative sample of X. There must be some relevant difference between X and Y to explain why such and such is true of X but not of Y.

Sometimes, reasoning with samples will involve a claim about some percentage or portion of the sample and target groups. Adding a **quantifier** will do this. A quantifier is simply a word that specifies an amount or quantity. There are many different kinds of quantifiers in English. We can identify a quantity in numerical terms, either as a fraction (e.g., one half or $\frac{1}{2}$) or as a percentage (e.g., 88 percent or 25 percent). But

we can also identify a quantity in non-numerical terms (e.g., most, almost all, a lot, a little.) Here is an example of such reasoning:

Such and such is true of 88 percent of X.
 X is a representative sample of Y.
 So, such and such is true of 88 percent of Y, too.

When reasoning with samples involves quantifiers, the reasoning is valid only if the quantifier in the conclusion is **appropriate** given the quantifier in the premise. In the following argument, the quantifiers are not appropriate:

Fifty-five percent of seniors surveyed claimed to have been in a car accident last year. The seniors in the survey were representative of seniors generally. It follows, then, that the vast majority of seniors are dangerous drivers.

Fifty-five percent is a majority; but it is not a *vast* majority. This argument exaggerates the findings from the survey. Notice also that, in this argument, the claim made about the sample group is not the same as the claim made about the target group. We saw in the previous section that the argument might not be valid if the claim made about the two analogues is not the same. In this case, the claim made about the sample group does not entail the claim made about the target group: being in a car accident is not the same as being a dangerous driver, and indeed being in a car accident is not sufficient for being a dangerous driver. This argument is thus invalid on two counts: its quantifiers are inappropriate, and the claims made about the analogues do not match up.

6.2.2 When Are Sample Groups Representative?

In assessing reasoning using samples, we need to know whether the sample group really is representative of the target group. The basic idea is clear enough: the sample group is representative of the target group just in case the two are alike in all relevant respects. But while this is clear, it provides little practical guidance in deciding whether some sample group is in fact representative of the target group.

6.2.2.1 Sample Size Let's start with an example that will help us bring out the main points. Suppose that a chef is planning to taste some samples of his soup in order to decide whether he should add more salt. How many spoonfuls should he taste before he has sampled enough

to know whether the soup needs more salt? This is a question about **sample size**: a sample is representative of the target group only if the sample is large enough. Just how big a sample has to be depends on how **homogeneous** the target group is with respect to the relevant property. Suppose that the soup in question is a simple clear chicken broth that we know has been well stirred. In that case, the soup is very homogeneous. That is, it is reasonable to expect in this case that the salt level in one part of the soup would be the very same as the salt level in any other parts of the soup. If so, then the chef would not need to collect a large sample—maybe even just one or two teaspoonfuls would be enough.

Let's consider a different example. Suppose that a biologist wanted to figure out the levels of arsenic in a local lake, and that the lake had several rivers feeding into it, and that some of them had lumber mills on their shores. In this case, the biologist probably would have to collect lots of samples, since it is likely that the chemical composition of the lake water is not very homogeneous from one part of the lake to another. The risk, of course, is that if the biologist took only one sample, it might turn out that that sample was not at all representative of most of the lake's water. The biologist might end up falsely concluding that the arsenic levels are fine in the lake as a whole. This shows that the required sample size depends on the homogeneity of the target group with respect to the relevant property: the more homogeneous it is, the smaller the sample has to be.

6.2.2.2 Random Samples The biologist example can also help us to see a second important point. Once she has decided to collect lots of water samples, how does she decide *where* to collect them? Obviously, collecting 100 samples from the very same spot would be a mistake. For if one sample of water at that spot is not representative of the water in the lake, then 100 samples from it wouldn't be representative either. This shows that while sample size is essential to its being representative, it is not by itself sufficient. Clearly, the biologist needs to collect water samples from as varied a set of locations as possible. She needs to collect water from the mouths of the various rivers, some from the very middle of the lake, etc. One way to ensure that her samples are representative of the lake water in general is to **randomly** select 100 sites on the lake from which to draw samples. (Let us assume that 100 samples would be enough. It might not be, but let's continue with this assumption, for the sake of the argument.) The idea of a random sample is a technical one: a sample of Xs is **random** just in case the probability of any one X being in the sample is the very same as the

probability of any other X being in the sample. A selection of jellybeans out of a jar is a random selection just in case the chances of any one jellybean being selected are the same as the chances of any other bean's being selected. Randomly selected samples are really valuable because a randomly selected sample that is large enough is *guaranteed* to be representative.

BOX 6.2.2A PRACTICAL TIP: COLLECTING RANDOM SAMPLES

Definition: A sample of Y is a random sample just in case every part of Y had the same chance of being in that sample.

Randomly selecting the sample will also protect against one source of possible research bias.

Of course, just because a sample is randomly selected, it does not follow that the sample is representative. To be representative, a sample also has to be of the right size. A single, randomly selected sample of the lake water will not be representative of the lake's water, if the lake's water is not homogenous from one location to another. Likewise, asking one randomly selected American for their views on global warming couldn't possibly be enough to draw a conclusion about what Americans in general think about global warming. Remember, to be representative the sample has to be just like the target in all relevant respects. If 10 percent of the target are working-class people, and if being working class is relevant to one's opinion on global warming, then 10 percent of the sample should be working class too. If 15 percent of the lake's water feeds in from one river, then 15 percent of the sample should be from that river too. The only way to make sure that the sample is like the target in the relevant respects is to make sure that the sample is large enough. That is why a sample has to be large enough in order for it to be representative.

6.2.2.3 Self-Selected Samples Sometimes, a sample group will be **self-selected**. A sample is self-selected just in case whether something is in the sample depends on whether he or she decided to be in the sample. This sometimes happens with those "instant" surveys on Internet sites, when anyone visiting the site is asked to answer a series of questions. The sample will be a self-selected one because the only people who answered the survey are people who decided to participate.

They selected themselves, as opposed to having been selected by the surveyors. Sometimes magazines have a questionnaire for their readers to complete and mail-in. The questionnaires that are mailed in will constitute a self-selected sample because the readers themselves chose whether to fill it in. Student evaluations of their classes will involve self-selected samples if the students can decide on their own whether to complete the survey. A self-selected sample is thus not a random sample. Still, a self-selected sample *might* be representative, but there is no guarantee that it is. It would just be a lucky accident that it is. For this reason, a self-selected sample group should be treated as nonrepresentative, unless there is independent reason to think otherwise.

BOX 6.2.2B SUMMARY: REPRESENTATIVE SAMPLES

1. To be representative, a sample group has to be of the right size.
2. Adequate sample size depends on how homogeneous the target group is with respect to the relevant property. The more homogeneous it is, the smaller the sample size has to be; the less homogeneous, the bigger it needs to be.
3. Randomly selected samples of adequate size are guaranteed to be representative.
4. Self-selected samples should be treated as nonrepresentative.

Of course, we might not know whether a sample is representative or not. As we have already seen, it is not always easy to know what the relevant respects are. And even when we do know this, it is not easy to know just how big a sample has to be. And even if we do know all of this, it is still not easy to tell whether a sample really was randomly selected. Polling companies usually include a statement of the margin of error of their poll, something like: “this poll is accurate to plus or minus 3 points, 24 times out of 25.” This margin of error is meant to measure just how certain the pollsters are about whether the sample really is representative.

BOX 6.2.2C CRITICAL THINKING IN RELIGIOUS STUDIES

A 2008 survey by the Pew Research Center reportedly found that 21 percent of self-proclaimed atheists in the United States believe in the existence of a God or Universal Spirit, and that one third of

these are fully certain of the existence of such a being. This is a surprising discovery, since being an atheist means *not* believing in the existence of a God or supernatural being. Something obviously went wrong in the survey. We can represent the reasoning as follows:

1. Twenty-one percent of the sampled atheists believe in the existence of God.
2. The sample is representative of American atheists.
3. So, 21 percent of American atheists believe in the existence of God.

The argument is valid, so because the conclusion has got to be false, at least one of the premises must be false.

Perhaps Premise 1 was false. That is, maybe the survey was not a reliable measure of the opinions of those who took it. Maybe it used a question that was ambiguous, biased, loaded, or otherwise bad. Or maybe the people who declared themselves atheists did not really understand the meaning of the word “atheist” or do not fully understand their own religious beliefs. Or perhaps they were somehow biased, and not being sincere either in claiming to be atheists or in reporting their religious beliefs their answers. Any of these possibilities might explain how Premise 1 might be false.

Perhaps Premise 2 is false. That is, perhaps the sample was not representative of American atheists. This would be so if the sample were too small. Maybe it was not a random sample. Since the survey asked the respondents to self-identify as atheists, this means that there was room for subject error and so for self-selection to have an effect.

Whatever the explanation, something went wrong in the survey, and using the critical thinking tools we have learned, we are able to describe what kind of error it might have been.

All of this means that, as critical thinkers, we need to be cautious before we accept the analogical premise of an argument by analogy that involves sampling. But being a critical thinker also means knowing what questions to ask, even if we cannot always find the answers

1. Is the sample large enough? How homogeneous is the target in the relevant respects?
2. Was the sample randomly selected?

If we know what questions to ask, and if we have the needed vocabulary, then we are ready to think reflectively about reasoning with samples.

BOX 6.2.2D MISTAKES TO AVOID: HASTY GENERALIZATION

It is a mistake to rely on an unrepresentative sample when reasoning using samples. This is a special case of the mistake of a false analogy, and involves a false or unacceptable analogical premise. (Keep in mind that an argument that makes this mistake might still have a true conclusion.)

The mistake is called a “hasty generalization” because the conclusion of the argument is a claim about the target group in general, and it is hasty because not enough care was taken to ensure that the sample was representative.

If you are constructing an argument using samples, then you have an obligation to ensure that the sample is representative.

If you believe that a sample is not representative, then you have an obligation to explain why.

6.2.3 Reasoning with Models and Maps

We have been studying how to reason using analogies, and have looked at the case of reasoning with samples. Reasoning by analogy is also very useful in cases where the phenomenon we are investigating is very complex, or very large, and it is easier to think about the phenomenon by thinking about a model or map of it. You may remember from grade school reasoning about the solar system by studying a model of it, with the planets and some of their moons, all orbiting the Sun. Biologists and geneticists do tests on animals in order to find out how various treatments would work on humans. Meteorologists and economists use computer models of storms or financial markets in order to understand and predict real storms and markets. We all use maps to find our way in a strange city. All of these are examples of reasoning about one thing by thinking about something that is a model or map of it. Evaluating this kind of reasoning is just the same as evaluating reasoning by analogy.

Let’s use an example to bring out the key points. Suppose that an economist wants to predict what would happen to housing prices in a certain region if interest rates were to rise and the level of unemployment was to fall. The economist’s model is dynamic in the

sense that she can affect it by changing the interest and unemployment rate. In designing the model, the economist used what she knows about housing markets and about the relations between them and interest and unemployment rates. She knows that some features of the market are irrelevant—such as the color of the house—and that others are relevant—such as the age or size of the house. Once her computer model is up and running, she can use it to predict how changes in real interest rates might affect the regional housing market. She could input a rise in interest rates of $\frac{1}{4}$ of a percentage point and then see how the computer model of the regional housing economy is affected. She might see lower sales or a difference in the average sale price. If the model is sophisticated enough, she might be able to see different changes in different parts of the regional market. The better the model represents the real market, the more reliable her predictions will be. In other words, the more the computer model is like the real model in all the relevant respects, the more reliable her predictions will be.

In designing her computer model, the economist will ignore features of the real situation that she believes are not relevant. We already mentioned that she might ignore the colors of the houses in the region, since it is unlikely that a house's color makes much of a difference to housing prices. As we have already seen, it is not always easy to tell what features are relevant to a given property and which ones are not. Our economist might discover that something that she thought was irrelevant is in fact quite relevant. If her model's predictions are always wrong, then one possibility is that her model is not a very good representation of the real market because it ignores something that is actually quite relevant. Indeed, this is one way that we can figure out what features of a situation are relevant and which are not. In other words, the process of trial and error involved in trying to construct a genuinely representative model of some phenomena can itself teach us a good deal about the nature of that phenomenon.

In designing her model of the regional housing market, the economist might deliberately ignore something that she knows is relevant, but only in very minor ways. Suppose that the number of windows in a house makes only a very insignificant difference to its sale price, and suppose that it would take a long time to collect and include window information in the model. She might decide to go ahead with the model anyway, leaving out the information about the number of windows. Her model would then be an **idealization**, pretending that the housing market is simpler than it really is. Strictly speaking, this means that her model is not an accurate representation of the real market. And it follows from this that it is not a perfectly reliable source of information on the real market. Still, this sort of idealization is acceptable, so long

as one keeps track of it. So long as she keeps in mind that the model's predictions are based on an idealization, and so long as she is right about how insignificant those idealizations are, she should be able to rely on it to provide fairly accurate predictions.

EXERCISE 6.2

1. Short-answer questions:

- a. What makes reasoning using samples a kind of reasoning by analogy? Use examples to develop your answer.
 - b. What is a representative sample? Use the four-step definition method from Chapter 2 in developing your answer.
 - c. Is a large sample guaranteed to be representative? Explain using examples.
 - d. Is a randomly collected sample guaranteed to be representative? Explain using an example.
 - e. Could a sample be representative even when it is not randomly collected? Explain using an example.
 - f. What is it for a target group to be homogeneous? Use the four-step definition method to develop your answer.
 - g. What is a self-selected sample and why do they tend not to be representative?
 - h. Suppose that Jones and Henry are both collecting a sample of Americans for a survey. Suppose that Jones' sample is randomly collected and that Henry's is self-selected. Is it possible that the two samples contain the very same members? Explain.
 - i. What is a hasty generalization, and what makes it a critical thinking mistake?
2. In assessing whether a sample is large enough, we need to know how homogeneous the target group is with respect to the relevant property. To your knowledge, are the following target groups very, somewhat, or not at all homogeneous with respect to the following properties:
- a. Males in your course; study habits
 - b. Males in your course; political opinions
 - c. Males in your course; anatomy
 - d. Mice; anatomy
 - e. Roses; genetic makeup
 - f. Soft drinks; sugar content

- g. Cars; safety features
 - h. Cars; reliability records
3. For the following groups, identify three properties with respect to which they are homogeneous and three with respect to which they are not:
- a. Maple trees
 - b. Planets
 - c. Sweaters
 - d. Cars
 - e. Humans
 - f. Shoes
 - g. TV shows
4. For the following pairs of samples and targets, identify one property with respect to which the sample is representative of the target, and one property with respect to which it is not:
- a. Ten people waiting for a bus; users of the city's public transit system
 - b. Ten people waiting for a bus; residents of the city
 - c. Fifteen hundred Americans responding to an online survey; Americans in general
 - d. Three mice; all mice
 - e. Three mice; humans
 - f. Five randomly selected cars; all cars
 - g. Five randomly selected cars; manufactured products
 - h. U.S. companies; worldwide companies
 - i. Three glasses of water taken at 8-hour intervals in your kitchen; the water in your pipes

BOX 6.3 CHAPTER SUMMARY

Reasoning by analogy is reasoning about one thing by comparing it to another that is just like it in relevant respects. This reasoning is valid, but it is usually not easy to tell what the relevant features are or to tell whether the analogues really are similar in all of those respects. Reasoning using samples and models are examples of reasoning by analogy.

6.3 MISTAKES TO AVOID

False Analogy It is a mistake for an argument by analogy to include a false analogical premise. An analogical premise claims that the analogues are alike in all relevant respects. There are two ways to be mistaken in making an analogical claim. One might be mistaken about which respects are the relevant ones, or one might be mistaken about whether the analogues are alike in those respects. You should accuse an argument of committing a false analogy only if you can identify a relevant difference. An argument by analogy using samples that has a false or unacceptable analogical claim as its premise commits the mistake called “hasty generalization.”

Unrepresentative Sample It is a mistake to draw a conclusion about a target group based on an unrepresentative sample. A sample is representative of a target just in case if the sample and target are alike in all the relevant respects. There are several things to consider in assessing whether a sample is representative:

- i. Is it too small? To be representative, a sample has to be large enough. How large it has to be depends on how homogeneous the target group is with respect to the relevant property. The more homogeneous it is, the smaller the sample size can be.
- ii. Was it randomly selected? A sample is randomly selected just in case every member of the target group has the same chance of being in the sample group. Randomly selected sample of the right size are guaranteed to be representative.
- iii. Was it self-selected? A sample is self-selected just in case the members of the sample are in the sample because they decided to be in it. A self-selected sample of the right size might be representative, but it is not very likely. It is best to treat self-selected samples as non-representative.

Fallacy of Idealization It is a mistake to ignore or overlook the respects in which a model or map of some phenomenon involves idealization. A model involves idealization when relevant features of the phenomena are deliberately ignored. This can be fine, but only so long as one keeps tracks of the idealization.

6.4 FROM THEORY TO PRACTICE: APPLYING WHAT WE HAVE LEARNED

Thinking Critically in Your Own Life We have been emphasizing that we can and should think critically in every aspect of our lives. This includes in our thinking about our own life, about what we want from life, about what kind of person we might want to be. In Chapter 1, you identified some features that you think are essential to being a morally good person and you were asked to give some reasons for thinking that they are in fact essential. In Chapter 2, you worked to construct definitions of them using the four-step definition method. Now that we have discussed what it is for reasons to be good ones—they must be acceptable and sufficient—try to construct an argument in a paragraph or two giving your reasons in such a way that it is clear that they are both acceptable and sufficient. You could use reasoning by alternatives, or about necessary and sufficient conditions, or by analogy. Ensure that the arguments that you construct are valid. Once you have constructed your arguments, identify the main weak points in them.

Thinking Critically in the Classroom This exercise is designed to help you identify the different forms of reasoning that you have to rely on in studying or engaging in your chosen discipline. We have studied several forms of reasoning: reasoning about alternatives, reasoning about necessary and sufficient conditions (including both reasoning about definitions, and reasoning about causal conditions), and reasoning with analogies (including the use of samples and models.)

- i. Look through one of the textbooks for your discipline, and find two examples of each of the kinds of reasoning we have studied in chapters 5 and 6.
- ii. Write out the reasoning in the form of an argument, making sure that it is valid, and assess the acceptability of the premises.

Thinking Critically at Work This exercise is designed to help you think critically about the kinds of reasoning you or the organization you work for rely on for success. In Chapter 1, you compiled a list of tasks that you regularly perform at work that require critical thinking. Look over that list, and identify those tasks that require you to reason about alternatives, or about necessary and sufficient conditions (including reasoning about definitions or about causal conditions), and those that require you to reason using analogies (including samples and models).

7

CRITICAL THINKING IN ACTION

Critical thinking is reasonable and reflective thinking aimed at deciding what to believe and what to do. In Chapter 1, we saw that part of what makes critical thinking reasonable thinking is that it requires that we have reasons for our beliefs. More specifically, we saw that thinking critically requires having epistemic reasons: reasons for thinking that our belief is true, or for accepting some claim that we are considering. In Chapter 2, we studied the ways that concepts and terms help to frame our investigations and our problems and we identified some practical strategies for clarifying and defining concepts and claims. Chapter 3 focused on what it is for reasons to provide **sufficient** support for a belief, and in Chapter 4 we studied what it is for our reasons themselves to be **acceptable** and, in particular, when we are justified in trusting the information we receive from observation, testimony, and measurement. In Chapters 5 and 6 we looked at several very common and very powerful **forms of reasoning**—reasoning about alternatives, reasoning about necessary and sufficient conditions, and reasoning with analogies. This final chapter is about pulling together the ideas, concepts, tips, and tricks we have learned into some practical strategies for helping us put critical thinking to work.

The aim of this chapter is to identify some general practical strategies that can help us to think critically at home, in our studies, and in

the workplace. The approaches are the same whether we are thinking about our own lives, about the discipline we are studying or engaging in, or about our tasks and responsibilities at work. Since you are probably reading this book in a college or university course, we will study how to apply the approaches in our studies. And since your discipline is likely not the same as most of the people in your class, we'll explore how these approaches can be applied across the curriculum. The point is not that what it is to think critically varies from one discipline to another; it does not. But disciplines do differ in the concepts they employ, in the sources of evidence they rely on, and in the kinds of reasoning that predominate. Remember that critical thinking is **reflective** thinking, and part of what this means is that thinking critically requires reflecting on the concepts one is thinking with, on kinds of evidence one is relying on, and on the kinds reasoning one is employing. One goal of this chapter is to provide you with some tools to help you be as reflective as you can.

7.1 THINKING CRITICALLY ABOUT A DISCIPLINE

There are three things involved in mastering a discipline. First, you need to master its **key concepts**. What are the key ideas, concepts, and terms that experts in the discipline use to state their claims, to formulate their hypotheses, and to analyze their evidence and data? Second, you need to master its **sources of evidence**. How do experts collect the data and information they need to answer their questions and to solve their problems? Finally, you need to master the discipline's primary **modes of reasoning**. How do experts in the discipline draw conclusions from the information they collect? In this section, we'll look at these three tasks in turn.

7.1.1 Identifying a Discipline's Key Concepts

In Chapter 2, we saw how key concepts are used to frame both problems and potential solutions. Part of what distinguishes one discipline from another are the concepts that experts in the discipline use to think about a phenomena. As we saw, different disciplines might study the very same phenomena from different points of view or perspectives, approaching they very same facts or puzzles with different conceptual resources, methods, and explanatory models. Indeed, this is part of why it can be so fascinating to study what experts in different fields have to say about some phenomena.

Consider the study of human sexuality. Sociologists and psychologists have a lot to say about it. But so do biologists and novelists. Even philosophers have tried their hand at making sense of it. To some extent, these researchers are interested in different aspects or elements of human sexuality. A psychologist might be more interested in exploring the central place that sexuality plays in our own self-conceptions, while a biologist might be more interested in the anatomical facts about how it works. But they might both be interested in the complex and subtle ways that sexuality interacts with social and group relations. The question: “Why do so many humans mate for life” will be understood and approached in very different ways in different disciplines.

Thinking about a phenomenon from the perspective of a given discipline requires using that discipline’s **key terms and concepts**. Thinking about human sexuality from a psychological perspective requires thinking about it using terms, concepts, and ideas that are not the ones that a biologist or a sociologist would use. Or, to change the example, the sport of football could be studied from many different perspectives. The sorts of questions and descriptions that a sociologist might give of a football game—perhaps focusing on the complex relations between individual players and the team, or between the fans and the team—would be very different from the kinds of questions and descriptions a sports physiologist would give—focusing on the anatomical features of players in different positions and the importance of different muscle groups. If, however, one wanted to think of it from a football fan’s point of view, one would have to use the concepts from the football rulebook. Different disciplines approach the very same phenomenon using different concepts and questions. Mastering a discipline requires knowing what its key terms and concepts are.

Identifying a discipline’s key concepts is not usually a very difficult task. An introductory textbook in the field will usually include a glossary of key terms. The textbook’s author will include in that glossary the words and concepts that she thinks one has to be familiar with to understand and engage in that discipline. Consider the key concepts and ideas in critical thinking. If we had to make a list of all of the technical terms we have been using in this book, we might produce the following list: epistemic reasons, nihilism, realism and skepticism, validity, argument, premise and conclusion, subargument, independent premise, truth, acceptable reasons, sufficient reasons, observation, testimony, reliable, measurement, accuracy and precision, necessary and sufficient conditions, exclusive disjunction, modus ponens, modus tollens, analogues. If one could master the meanings of each of these words (especially the word “validity”!), then one would be well on the

way to mastering the study of critical thinking itself. (This, however, would not be the same thing as being a strong critical thinker. I might be an expert on football, but a terrible football player!)

7.1.2 Clarifying a Discipline's Key Concepts

Of course, mastering a discipline's key terms requires more than just being able to make a list of them. One has to know how to use them properly and for this it helps to be able to say or explain what they mean. In Chapter 2, we studied a method for constructing and evaluating definitions. We can use that method to define a discipline's key terms. Here is an example from the study of critical thinking:

An argument is valid just in case if its premises were true, then its conclusion would have to be true too. In other words, it is not possible for the premises of a valid argument to be true and yet for its conclusion to be false. There is no way that its conclusion could somehow turn out to be false if its premises are in fact true. The premises provide conclusive evidence for the conclusion. For example, the following is a valid argument: All men are mortal; Socrates is a man; so he is mortal too. Here is an example of an invalid one: All mortals are men; Socrates is a man, so he is mortal too. The first one is valid because if its premises were true, then it would have to be true that Socrates was mortal too. But the second argument is not valid: maybe Socrates is a mortal cat! A valid argument is not the same as an argument that has true premises or that has a true conclusion. Whether an argument is valid is a matter of whether the conclusion would *have to be true, if the premises were true too*. An argument that is both valid and has true premises is called a sound argument. It is important not to confuse validity with soundness.

In this explanation of the meaning of the word "valid" the first sentence contains the initial slogan-like version of the definition. It is useful to have a short and simple to remember statement of the meaning. The next three sentences provide a longer expansion or elaboration of the slogan, putting the idea in other, equivalent terms. Then there are a few examples. Examples may not be useful or convenient for some sorts of definitions. Then the text provides some contrasting concepts or terms. Since the goal of providing a definition is to help to avoid confusions, it is very useful to contrast the concept being defined with nearby concepts, ones that someone might easily mistake or confuse with the one being defined. This four-step method is not the only possible way to provide a definition. But it is extremely handy, and as we

saw in Chapter 2, it can be used to explain what is meant by a statement or claim as we all to explain the meaning of some concept or term.

Let me repeat something from Chapter 2. The definition of validity I just gave provides necessary and sufficient conditions: it states what all and only valid arguments have in common, and it does this in a way that provides a rule to tell for any argument whether it is valid or not. But this is not always possible, and sometimes it is not even very desirable. The key terms that define a discipline are usually not so precise that they can be given a definition in terms of conditions that are necessary and sufficient. In practice, this allows room for flexibility. It allows researchers to formulate hypotheses, raise questions, consider objections, and state alternatives that might not be possible if each and every word had a precise definition.

BOX 7.1.2 SUMMARY: IDENTIFYING AND DEFINING A DISCIPLINE'S KEY CONCEPTS

Mastering a discipline requires mastering its key concepts. These can usually be found in introductory textbooks. But knowing what they are is not enough. One has to be able to use them to think about the phenomena, and for this it is useful to use the four-step definitional method discussed in Chapter 2.

EXERCISE 7.1.2

1. Visit the library and find an introductory textbook on human sexuality from psychology, sociology, evolutionary biology, and chemistry. Compare and contrast the textbooks' key words.
2. For your own discipline: identify the ten most important concepts and terms. Using the four-step method, compose a definition of each one.

7.1.3 Identifying a Discipline's Sources of Evidence

We have been discussing the way that mastering a discipline requires mastering its key concepts. But mastering a discipline also requires understanding what kinds of evidence it relies on, and this is usually a matter of knowing what kind of sources it uses. In Chapter 4 we explored three kinds of sources: observation, testimony, and measurement. To some extent, every discipline relies on all three sources of evidence.

Indeed, it is hard to imagine any serious branch of science—pure mathematics and philosophy aside—that do not depend on observation, testimony, and measurement. As we know, part of what makes critical thinking reflective thinking is that it involves making explicit the sources of evidence that we are relying on as we go about deciding what to believe or what to do. Let's consider a couple of examples.

Suppose that a criminologist wanted to find the most cost-efficient form of incarceration. What sorts of information would he need? He might visit several institutions, doing direct observation of the conditions at each one. He might interview managers and government officials, seeking their views on management practices and on the budgets and financial conditions of their institutions. He might develop a questionnaire to learn about the inmate's attitudes. In other words, he would rely on evidence from direct observation (the on-site visits), testimony (the interviews), and measurement (the survey). This study, in other words, might well require relying on evidence from all of the kinds of sources of evidence that we have studied in this book.

Sometimes, a discipline will rely on a source of evidence that is unique to it. This is the case in the brain sciences, where researchers use highly sophisticated brain imaging techniques. While the basic technology that produces these images is used in other disciplines too, it has been developed specifically for use in studying the brain. It provides information that simply cannot be collected in any other way. Indeed, until fairly recently, scientists studying the human brain had very little information to go on, since our access to living active brains was extremely limited. But now with the invention of brain imaging techniques, we are able to collect an enormous amount of information about how the human brain changes and grows, about its structure and functional organization, and about how various diseases and accidents can affect it. The development of brain imaging technology provided a new source of evidence and information, one that helped the scientific study of the brain make greater progress than it had for hundreds of years.

BOX 7.1.3 MISTAKES TO AVOID: APPEALING TO IGNORANCE

As we saw in Chapter 5, there are many cases where measurement is difficult. In the case of deciding what to do, it can be difficult to compare the costs and benefits of different proposals if the costs and benefits are unpredictable or if they are incommensurable.

But it is a mistake to ignore potential costs or benefits of a proposal on the grounds that the costs or benefits are either unknowable or incommensurable. Doing so would in effect be appealing to ignorance.

EXERCISE 7.1.3

1. What sources of evidence do the following disciplines rely on? Be as specific as you can.
 - a. History
 - b. Nutrition
 - c. Political science
 - d. Early childhood education
 - e. Social work
 - f. Architecture
 - g. Restaurant reviewing

2. What kinds of evidence might be needed to do the following studies? Be as specific as you can.
 - a. To compare the effectiveness of relaxation over medicine for treating ordinary headache pain.
 - b. To identify the impact on commuting of adding a stop sign at a busy intersection.
 - c. To discover whether there is a greatest prime number.
 - d. To understand the influence of popular music on fashion styles.
 - e. To complete a review of a new Broadway show.
 - f. To learn the effect that taking a critical thinking course has on a student's subsequent university education.
 - g. To find which of three paint varieties is most mildew resistant.
 - h. To learn the effects of substituting baking soda for baking powder in a cookie recipe.
 - i. To learn the impact that stay-at-home fathers have on their children's socialization.

7.1.4 Identifying a Discipline's Modes of Reasoning

Finally, mastering a discipline requires mastering the kinds of reasoning that experts in it rely on. In this book, we have studied three main forms of reasoning: reasoning about alternatives; reasoning about necessary

and sufficient conditions (including causal conditions); and reasoning with analogies (including sampling and modeling). As with the sources of evidence, most disciplines use all of these modes of reasoning to some extent or another.

Consider again our criminologist trying to identify the most cost-effective prison system. What kinds of reasoning will he engage in once he has collected his information and data? He might reason about alternatives, for he might have started with several candidates for being the most-efficient system, and is hoping that his data and evidence will help him narrow down the list to one or a small number of equally efficient systems. He might reason about necessary and sufficient conditions too, since he would need to begin with some definition or criterion of "most efficient." Indeed, he might combine these two forms of reasoning, ruling out certain candidates by showing that they lack features that are necessary for satisfying that definition or criteria. Since the researcher is interested in comparing and contrasting different prison systems, there will have to be some reasoning by analogy as well. This study will thus employ all of the forms of reasoning we have discussed in this book.

We have been discussing how the researcher might use these three forms of reasoning to draw conclusions from the evidence of data he has collected. But these forms of reasoning are often employed at other stages in an investigation as well. We have already noted that the researcher will likely begin with some sort of conception of what it is for a prison to be cost-efficient. Developing and checking this conception will likely involve reasoning about necessary and sufficient conditions, and may involve using the counter-example method to help identify just what makes a prison system efficient. Indeed, it is likely that the conception might change and improve as the researcher begins to collect the data: it is only in very unusual cases that the definitions are crystal clear at the beginning of an investigation. Designing the questionnaire to uncover inmate opinions also requires reasoning: one has to decide what questions to ask and ensure that they are well formed and not biased or leading. Identifying a sample group requires thinking about the prison population in general in order to ensure that the sample is representative.

Reasoning about alternatives, about necessary and sufficient conditions, and about analogues can thus occur at many different stages in an investigation. It might occur at the very beginning as the researcher works to formulate and clarify her hypotheses. It might be involved as she decides on what sources of evidence to use and on what methods of measurement to employ. It will guide her as she

makes those measurements and collects the data. And it will be her guide as she draws conclusions from the evidence she has collected. Indeed, this kind of reasoning is pervasive as we try to decide what to believe.

7.2 CRITICAL THINKING QUESTIONS

So far in this chapter, we have been studying three things one needs to learn in order to master a discipline. One must learn the discipline's **key concepts**, its **sources of evidence**, and its **modes of reasoning**. Knowing what kinds of questions to ask is of great value in thinking critically. Throughout this book, we have looked, sometimes in great detail, at different elements of critical thinking. But the kinds of questions we have learned to ask can helpfully be categorized into three kinds: questions about meaning; questions about truth; and questions about value.

One can ask questions about **meaning**:

- i. What is the claim being defended, or the course of action being proposed?
- ii. What are the claim's key words and what do they mean?
- iii. Can we provide examples to illustrate them and can we identify some contrasting concepts?
- iv. What is the framework within which this claim or proposal is being raised?
- v. What are some alternative or contrasting frameworks?

One can ask questions about **truth**:

- i. Some questions about truth are about the **acceptability** of premises. What are the sources of the information in the premises? Are those sources reliable? What objections to the truth of those premises are there, and how are they to be addressed?
- ii. Some questions about truth are about the **sufficiency** of premises. What forms of reasoning are in use? Would those premises constitute a valid argument for the conclusion? What conclusions could be validly drawn from those premises?
- iii. Some questions about truth are about **alternatives**. What other views on this subject are there? What are the strongest reasons in their favor? What are the strongest objections to them?

One can ask questions about **value**:

- i. Why is it important to perform this study?
- ii. What is the context that makes this study important or interesting?
- iii. Would the answers to these questions impact or influence studies in other disciplines?
- iv. How does this conclusion fit into the author's broader argument?

BOX 7.2 PRACTICAL TIP: THE RULE OF THREES

An important skill in thinking critically is being able to construct alternatives and to identify contrasting concepts. It is a useful rule to try to find at least three alternatives or contrasting concepts. Often, finding the first two will be relatively easy—it is finding the next one or two that proves enlightening.

Critical thinking is aimed at deciding what to believe or do. We think critically whenever we try to decide what some evidence shows about some question or phenomena, whenever we assess or evaluate what someone else believes or is arguing, and whenever we try to present our own reasons for believing something. Keeping these three kinds of questions in mind can help us to organize our thinking. As we read or listen to someone develop their reasons for believing something, we should be asking these three kinds of questions. As we plan our own investigation we should keep in mind the importance of defining the key concepts and problems clearly, of ensuring that our sources are reliable and trustworthy, and of keeping an eye on the big picture.

7.3 THINKING CRITICALLY IN YOUR OWN DECISION MAKING

Sometime during your university career, you will be expected to write an essay or report in which you present and defend a point of view. Indeed, it is quite likely that you already do this on a regular basis. And it is inevitable that once you start a career you will be expected to present suggestions and recommendations. What we have learned in this book can help you think about how to structure your essay or report.

7.3.1 Clarify Your Views

The single most important thing is to be as clear as you possibly can about what your claim or proposal is. If your reader is not clear what exactly you are claiming or what course of action you are recommending, then there is little chance that you will persuade them, and some chance that any agreement will be based on a misunderstanding. If your claim or proposal is formulated using technical or specialized concepts, make sure that you carefully and clearly define them. We have seen, in Chapter 2, that the four-step definitional method can be used to state a claim or proposal clearly. In many cases, the most important thing is to prevent possible misunderstandings by explicitly noting views or proposals that you are not recommending but that you suspect others might confuse with the one you are recommending. Doing this is also a useful way for you to make sure that you are clear in your own mind about what your views are.

Usually, the claim or proposal that you are advancing will be part of a much larger investigation or project, and it is often very helpful to make explicit what this larger context is and how your piece fits into it. This will help your audience to locate your claim or proposal in a larger picture, and will help to make its importance or value of your position clear.

7.3.2 Clarify Your Reasons

It is very important not to confuse your view or proposal with your **reasons** for accepting that claim or proposal. A statement of a view or of a possible course of action might take an entire paragraph, especially if you are using something like the four-step method. Stating the reasons for your view might take as long. Sometimes, it might be helpful to formulate your reasons in the form of an argument, and then spend a paragraph clarifying each premise and the conclusion. There are different ways to organize your reasons, but none will succeed if the reader is not able to easily tell the difference between your reasons and your view.

There are many reasons for keeping a sharp line between one's views and one's reasons. First, someone might agree with your view or proposal but for different reasons. They might agree with you that building a bridge is the best option, but not for the same reasons as you. Perhaps you think it is the most cost-effective option and they think it is the option that will do most to resolve traffic congestion. Second, someone who resists your view might in fact accept your reasons, and might not

have realized that they support your view. By laying out the reasons clearly, and showing that they do in fact support your view, you could bring them around to your side. Third, someone might be able to offer additional (or even better) reasons in support of your view.

7.3.3 Show That Your Reasons Are Acceptable and Sufficient

No matter what your view is or what your reasons are for it, there are bound to be objections, and it is important to be as honest as you can about what they are and about how to respond to them. But as we know, there are really only two kinds of objections. Someone might object that your reasons are not **sufficient** to support your claim or proposal, or they might object that your reasons are not **acceptable**. (Or both!) The fact that people will raise objections to your view or to your reasons is actually a good thing, since the goal of critical thinking is to get at the truth. Objections are simply a healthy way of making sure that your reasons are sufficient and acceptable.

The best defense, they say, is a strong offense. One way to go on the offensive is to make it as clear as you can in your essay or report that your reasons are sufficient and acceptable. We have seen several forms of reasoning that are guaranteed to be valid (and so sufficient) and it would be good for you to try to present your reasoning in a valid form. Doing this is relatively easy if your reasoning is about necessary and sufficient conditions or about alternatives. But it is more complicated if you are reasoning about causal relations or using analogies. But the strategies identified in chapters 5 and 6 can be helpful.

Objections to the acceptability of your reasons are more serious, and more difficult to defend against. It is good to be as clear as you can about the source of your reasons. If you are relying on observation, say so and do what you can to show that your observations were collected in optimal conditions. If you are relying on testimony, explain why the witness is appropriate, competent, and unbiased. If you are relying on measurements, say something about the accuracy and reliability of the measuring device. If you are reasoning about alternatives, make it clear that your disjunction is acceptable by showing that it is exhaustive. If you are reasoning about necessary and sufficient conditions, discuss whether the conditional states a definition or a causal relation and do what you can to show why it is acceptable. Considering and responding to possible counter-examples is a useful strategy for this. If you are reasoning using analogies, discuss what the relevant respects are and the reasons for thinking that the analogues are alike in those respects.

We have learned enough in this book to know that these tasks can be difficult. The most important thing, though, is to be as honest as you can be about just how acceptable you think the reasons or pieces of evidence are. If you suspect that there are difficulties with it, say so. Given our incomplete information, we are bound to have to make decisions about what to believe and do on the basis of evidence whose acceptability we are not able to prove beyond a shadow of a doubt. While admitting that you are sensitive to the question of acceptability may not make your argument any stronger, it will show you to be a strong critical thinker, and so increase the chances that others will be swayed to your side. And it will help your case if you can discuss these issues using the key critical thinking concepts.

7.3.4 Identify and Respond to Alternatives

No matter how good your reasons are for your view or proposal, there are bound to be alternative views and proposals. Sometimes, people will hold the opposite view, and will have reasons that they think are sufficient and acceptable. Sometimes, though, people will hold a view that is close to your view but different in subtle ways, and will have reasons that they think are sufficient and acceptable. It is important to identify and respond to these alternatives. We have already seen that identifying them can be a useful thing to do as part of clarifying your own view and reasons.

Responding to alternatives involves doing two things. First, it is good to raise **direct objections** to the alternative itself. Direct objections are reasons to think that the claim is false or that the proposed course of action is not a good one, or not as good as the one you are proposing. Second, it is good to raise objections to the reasons that have been or might be offered in support of that alternative. If the alternative you are considering is to build a tunnel rather than a bridge, then those who support the tunnel option will probably have reasons on their side, and it is good for you to identify what they are (making them as clear as you can so as to avoid the straw-man mistake) and then explain why you think they are either not sufficient or not acceptable. Perhaps you think they rely on evidence from a biased witness, or are based on measurements from an unreliable device. Whatever your objections might be, it is important to be as clear as you can about what they are. And for this, there is nothing better than using the key critical thinking concepts to frame your objections.

BOX 7.3 SUMMARY: THINKING CRITICALLY IN YOUR OWN REASONING

In presenting your reasoning for some claim or proposal, it is important to:

- State your view clearly
- Separately state your reasons
- Defend your reasoning by discussing the acceptability and sufficiency of your reasons
- Consider alternatives and identify objections to them

7.4 THINKING CRITICALLY IN DISCUSSION

Thinking critically in a conversation involves knowing what kinds of questions to ask. But it also involves maintaining a kind of critical and emotional distance. And most crucially, it involves insisting on reasons.

7.4.1 Ask Open-Ended Clarification Questions

It is very important to ask for **clarification**. We are all familiar with the way that a disagreement that seemed at first to be quite substantive—about something very deep and important—can suddenly turn out to be a mere linguistic disagreement; a disagreement, not about the facts, but about how to state the facts. This is perhaps especially true in some of the most emotionally charged debates of our time. For this reason, it is important to make sure that you know exactly what people you are in a discussion with mean by their words. Of course, they have an obligation to be as clear as they can be, and to do whatever they can to avoid or prevent confusions and misunderstandings. But as a critical thinker, you have the same obligation. In a discussion, when you are the listener, you can fulfill this obligation by asking the right kinds of questions.

Questions that have a “yes” or “no” answer are usually not as informative as questions that require a long answer. This is especially true if an apparent disagreement rests on linguistic differences or misunderstanding. You want to do whatever you can to rule out this possibility and ask open-ended questions that require the speaker to restate her view in other ways, to elaborate or expand on her view, to offer examples of it, and to identify contrasting views is the way to do it. In other words, use the four-step method to help develop open-ended

questions. And you can ask these questions not just about her conclusion, but also about her reasons and premises.

7.4.2 Withhold Disagreement and Agreement

We all know what it feels like to explain our opinions and views to people we know disagree with us. It puts additional pressure on us, pressure that is not really helpful if the goal is to get clear on the truth of the matter. For this reason, it is best to withhold disagreement until the speaker has had a full opportunity to explain, defend, and support her view. If she feels rushed into explaining it, she might leave out some crucial distinctions, and this might have the effect of weakening what might in fact be a fairly plausible position. Or she might skip over some of the premises or fail to mention a response to an objection. Your goal as a critical thinker in a discussion is to help the other people make a case for their views that is as clear, as strong, and as complete as possible.

It may be just as important to withhold agreement as to withhold disagreement. After all, you and the speaker might agree on the truth of the conclusion for very different reasons. Perhaps you would reject his reasons as either unacceptable or insufficient. Again, your goal is to help the speaker develop as clear, as strong, and as complete a case for his view as possible, and if you and the speaker reach agreement too quickly, before all of the elements of his view have been laid out, your superficial agreement will end up hiding an underlying and perhaps very significant disagreement. It is, for this reason, just as important to question thoroughly people you agree with, as it is to question thoroughly people with whom you disagree.

BOX 7.4.3A MISTAKES TO AVOID: TAKING AGREEMENT FOR GRANTED

We saw in Chapter 2 that it is a mistake to distort or misrepresent someone's views or reasons. Not only is this a rather rude thing to do, but it is also an obstacle to finding the truth.

The strawman fallacy usually occurs when someone states or describes the opinions or reasons of someone with whom they disagree. But a related mistake is to assume that there is agreement on views or reasons without first checking. Asking open-ended questions while withholding both agreement and disagreement can help ensure that all parties know the common ground.

7.4.3 Keep Emotional Distance

Being a critical thinker means that our beliefs should be based on epistemic reasons, and not on emotional or pragmatic ones. Basing one's beliefs on emotions rather than on epistemic reasons is a mistake, since how a claim or proposal makes one feel emotionally is not a reliable guide to whether the claim is true or the proposal good. Emotions can also make it difficult to collect the evidence we need, or even from investigating further. Emotions get in the way when we identify too much with our own opinions and beliefs or with our own methods for collecting or evaluating evidence. If I become too emotionally attached to my beliefs and opinions, then I may react negatively when someone asks me for my reasons, or when they raise objections to my belief or when they state their own alternative beliefs. I might feel that they are criticizing me and not just my beliefs. The same is true if I am asked to defend my assessment of the evidence or my use of different methods for collecting evidence. If I come to identify too closely with these particular methods for assessing and collecting evidence, if I come to think of my value as a researcher as tied into their value, then I will react to criticisms of them as if they were criticisms of me and my judgment. This feeling of being under attack might make me feel defensive, and this can prevent me from thinking critically about the issue at hand. The same is true when I ask someone for his or her reasons. This sort of question is easily taken as aggressive or combative, even when the intention is simply to consider the issue from all sides as thoroughly as possible.

BOX 7.4.3B PRACTICAL TIP: DON'T PERSONALIZE REASONS

Reasons and evidence do not belong to anyone; in this sense they are **universal**. And whether they are good has nothing to do with who accepts them; in this sense they are **objective**. To avoid personalizing reasons, replace the following:

- a. What evidence do you have?
- b. What are your reasons?
- c. Why do you believe that?

with the following impersonal ones:

- a'. What evidence is there?
- b'. What reasons are there to believe that?
- c'. Why should we believe that?

Knowing how to **distance** oneself from one's beliefs and opinions in order to think critically about them is not easy. It is one of the hardest things to achieve. But the best way to avoid this feeling is making sure that one's beliefs and opinions are based on enough of the right kind of evidence. Again: **think twice; decide once**. Another strategy is to avoid talking about "my reasons" or "your reasons" and to talk instead of "the reasons" or "some reasons." This makes sense anyway, since reasons and evidence are not owned or possessed by anyone: they are universal and objective. Instead of asking "What are your reasons for believing that?" which can come across as confrontational, ask, "What reasons are there to believe that?" which makes the question sound less confrontational. Instead of asking, "What is your evidence?" you can ask, "What evidence is there for that?"

BOX 7.2.3C SUMMARY: THINKING CRITICALLY IN DISCUSSION

Since the goal of critical thinking is to get at the truth, it is helpful during a discussion to:

- *Ask open-ended clarification questions*, in order to allow everyone to fully state their views and reasons;
- *Withhold agreement and disagreement*, in order to find common ground and avoid merely linguistic disagreements;
- *Keep your emotional distance*, in order to enable a friendly and cooperative search for the truth.

7.5 FROM THEORY TO PRACTICE: APPLYING WHAT WE HAVE LEARNED

Thinking Critically in Your Own Life Throughout this book, we have been emphasizing that we can and should think critically in every aspect of our lives. This includes in our thinking about our own life, about what we want from life, about what kind of person we might

want to be. In Chapter 1, you identified some features that you think are essential to being a morally good person and you were asked to give some reasons for thinking that they are in fact essential. In Chapter 2, you worked to construct definitions of them using the four-step definition method. In Chapter 6, you constructed an argument in a paragraph or two giving your reasons. Now you are in a position to develop it into a full defense.

- Using the concepts that we have studied, discuss the sufficiency and acceptability of your reasons.
- Identify possible objections to your view, state them clearly and using the concepts that we have learned, respond to them.
- Identify some alternative views, state them clearly, present the very best reasons in support of them, and raise objections to them.

In writing it out, pretend that your audience is someone who has never taken this course or read this book.

APPENDIX A

MISTAKES TO AVOID

Critical thinking is reasonable and reflective thinking aimed at deciding what to believe and what to do. Throughout this book, we have identified mistakes that a good critical thinker should avoid. Some are mistakes that can arise in clarifying or defining a view. Others are mistakes that can arise as we collect or rely on evidence or reasons for a view. Still others arise when we try to draw conclusions for our evidence. And there are even mistakes that can arise as we assess other people's views or reasons. Knowing what they are will help us to avoid them in our own reasoning. But it will also help to make it clear just what the value is in being a critical thinker: thinking critically is valuable in part because it helps us to avoid some mistakes. This appendix lists all of the mistakes we have discussed.

Appeal to Origins It is a mistake to assume that a belief's originating reasons are **epistemic** reasons too. There are many factors that influence what one believes, and not all of them need be epistemic. Jones' belief that humans evolved from other species might have been originally based on excellent reasons he learned in school. In that case, the originating reasons would be epistemic ones. But maybe he took on that belief as an act of rebellion against his parents. In that case, his

original reasons for believing that are not at all epistemic ones. It is a mistake to assume that a belief's originating reasons are epistemic reasons too.

Personalizing Reasons It is a mistake to personalize reasons by treating them as if they belonged to someone. Epistemic reasons are **universal**: if they are reasons for me to believe something then they are equally reasons for anyone else to believe something. Reasons are also **objective**: whether they are good reasons has nothing to do with me or with anyone else. Personalizing reasons can obscure the fact that they are universal and objective. Identifying with one's views or one's reasons, or defining another person by their views or reasons, can allow emotions to get in the way of thinking critically and can make it difficult to achieve and sustain the level of thoughtful reflection that is essential to critical thinking.

Appeal to Relativism It is a mistake to assume that truth is relative. Relativism with respect to some subject matter is the view that the facts in that area are in some way dependent on our beliefs about them. We noted in Section 1.2.1 that relativism might be the right attitude to take towards such topics as what is humorous or what is tasty. But for most topics, even religious and moral ones, it is best to assume that realism is the appropriate attitude, unless one has powerful reasons not to. For most topics, in other words, it is wrong to assume that what is true for me might not be true for you, or that what is true for our community or culture might not be true for others. Truth is the same for everyone.

Sometimes, an appeal to relativism will be used as attempt to bring a discussion to an end. One person, perhaps tired of the debate or feeling that they are on the losing side, will say to the others: "Well, I'm entitled to my view and you are entitled to yours." This kind of response is fine if what is intended is that everyone is allowed to make up their own minds about what to believe or do. But if the point is that we can both be right even when we disagree, then this is a mistake that we should avoid, unless there is excellent reason to think otherwise.

Appeal to Emotion It is a mistake to base a belief only on emotions. For a belief to be justified enough for knowledge it must be based on sufficient and acceptable epistemic reasons. Epistemic reasons are reasons to think that the belief is true. Emotional reasons are not epistemic ones. How a belief makes us feel has nothing to do with whether

the belief is true. As we have already noted, critical thinking does not aim to eliminate emotion from our decision-making. I doubt this would be possible even if it were worthwhile. Many of our beliefs are so fundamental to our deepest conceptions of ourselves, of our culture or our place in the universe that the pain involved in abandoning them would be too great to bear. It is fine for our beliefs to have or even for them to contribute to such emotional supports, so long as they also have sufficient support from acceptable epistemic reasons. It is a mistake to base our beliefs on nothing but emotional reasons. It is also a mistake to allow emotions to prevent us from collecting or assessing the evidence we need to make the decisions we must.

Privileging Available Evidence It is a mistake to assume that evidence that we currently have is more acceptable or more sufficient than evidence that we might collect. It might be that our current evidence is the best we can get. But we will not know this without good reasons. Usually, this means trying to collect more. Crucially, even if we have excellent reason to rest content with the evidence we have, we should always keep an open mind that we might uncover new evidence that will over-ride or undermine the evidence we now have.

Appeal to Tradition It is a mistake to believe or do something simply because that belief or practice is traditional. The fact that a belief or practice has a long history is not an epistemic reason to continue it. Its history cannot show that the belief is true or that the practice is worthwhile. But as with appeals to emotion, the point is not that we should work to avoid relying on tradition. Indeed, we probably could not know everything we do if we did not rely on others. In Chapter 4, we considered when we should trust the evidence provided by other people including our ancestors. We saw that it is not that hard to decide when we have good epistemic reasons to believe what they tell us or to maintain the practices they pass down. But it is a mistake to rely on what others tell us without also relying on epistemic reasons.

False Definition It is a mistake for a definition to be too broad (that is, for it to include things that do not fall under the concept) or too narrow (that is, for it to exclude things that do fall under the concept defined), or both. A **counter-example** to a definition is an example that shows that the definition is too narrow or too broad. The four-step method can help us to avoid this mistake by requiring us to look for counterexamples and contrasting concepts.

Equivocating on the Meaning of “Wrong” It is a mistake to confuse the different meanings of the word “wrong.” We can distinguish at least three dimensions of evaluation that we might have in mind in calling an action wrong. The action might be legally wrong, in the sense that it violates a law or legal statute. The action might be prudentially wrong, in the sense that it has consequences that would be bad for your health or wealth. The action might be morally wrong, in the sense that it violates moral rules or sanctions. These three kinds of “wrongness” are quite different, and an action may be wrong in one sense but not another. Lying to your best friend might be morally wrong, even though it is not legally or prudentially wrong. In defending or evaluating reasons to believe that some proposal is wrong, it is very important to be as clear as one can about what kind of wrongness is at issue.

Talking Past Each Other It is a mistake for participants in a discussion not to recognize that they mean different things by the key words and phrases they use. This can be recognized and avoided by asking open-ended clarification questions, by allowing everyone the time they need to fully state their views and their reasons, and by a careful use of the four-step method.

Being Stuck in a Framework It is a mistake to frame or describe a problem in a way that prevents one from seeing certain alternatives or solutions. The four-step method can be helpful in avoiding this mistake, by requiring one to identify closely related but contrasting descriptions of the problem. Sometimes, reframing a problem by describing it in different terms or from a different disciplinary point of view can help resolve it. This is what “thinking outside of the box” requires.

Strawman It is wrong to distort or misrepresent another person’s views or reasons. Not only does this hurt their feelings (and when done deliberately is exceedingly boorish), but it undermines the goals of critical thinking, which is to seek the truth in another person’s views, even if those views are not our own. The four-step method can help us to avoid this by requiring us when defining another person’s proposal or belief to identify closely related but subtly different contrasting proposals or claims.

Confusing Undermining and Overriding Evidence It is a mistake to think that evidence from an unreliable source must be false. A belief is undermined when we have good reason to think that it is based on evidence from an unreliable source. A belief is overridden when we

have good reason to think the belief is false. When we have evidence that some source is unreliable, then we have good reason not to accept that evidence. But for all that, the belief might still be true. Indeed, evidence from an unreliable source might nonetheless be true. Just because the witness is not properly trained or is biased against the defendant, this does not mean that her testimony is mistaken.

Untrustworthy Testimony It is a mistake to accept testimony from a witness if you have reason to believe that the topic is inappropriate, or that the witness is incompetent or biased. Testimony is appropriate only on topics for which there are recognized experts. An expert must be properly trained and properly informed. And a witness must not be biased—motivated to lie about, to exaggerate, or to understate the facts. This mistake is sometimes called the fallacy of appeal to authority. But it is important to keep in mind that without acceptable testimonial evidence none of us would have very much knowledge at all. Most of what we know (or at least, what we think we know) we learned from others.

Confusing the Reliability of a Measuring Device with Its Precision or Consistency It is a mistake to confuse the reliability of a measuring instrument with its consistency or its precision. An instrument is reliable just in case it provides accurate information. An instrument is consistent just in case it provides the same measurement at different times. The precision of a measurement is a matter of how finely grained it is. A measuring instrument, such as a ruler, a gyroscope, or a questionnaire, might be reliable even though it is neither consistent nor precise; and neither consistency nor precision is sufficient for validity.

Confusing the Value and the Likelihood of a Cost or Benefit It is a mistake to confuse how good or bad some consequence would be with how likely it is. Winning the lottery would be terrifically good, but it is extremely unlikely. When comparing the costs and benefits of alternative proposals, you need to factor in the likelihood of those costs and benefits, as well as their value. A course of action that promises very high value but at very low probability (e.g., spending your money on the lottery), might not be as good as one that promises a high probability of moderate value (investing cautiously for the long run).

Bad Question It is a mistake to ask a question that is ambiguous, contains charged or slanted words, or that hides a controversial presupposition.

False Alternative It is a mistake to reason about alternatives with a false disjunction. Not only does this make the disjunction an unacceptable premise, but it also guarantees that the conclusion of the reasoning—even if that reasoning is valid—will be false.

Lucky Disjunction It is mistake to reason about alternatives with a disjunction that is not exhaustive. Such a disjunction might still be true—but one would not know it to be true. One’s reasoning would reach a true conclusion only by luck, and it is wrong to rely on luck over careful thought.

Affirming a Disjunct It is a mistake to conclude from the fact that one disjunct of a disjunction is true that the remaining disjuncts are false, unless one knows that the disjunction is exclusive. A disjunction is an either-or statement and is true just in case one of its disjuncts is true. A disjunction is exclusive just in case *at most* one of its disjuncts can be true (that is, each disjunct excludes the other). In the case of a nonexclusive disjunction, both disjuncts could be true. Affirming a disjunct is reasoning in the following way:

Either P or Q
 P is true.
 So, Q is not true.

This reasoning is valid only if the disjunction is an exclusive one.

Denying the Antecedent It is a mistake to conclude that the consequent of a conditional is false just because the antecedent is false. A conditional is an “if, then” statement, and claims that the truth of the antecedent (the part after the “if”, would be sufficient for truth of the consequent (the part following the “then”). Denying an antecedent is reasoning as follows: If P, then Q; it is not the case that P; so, it is not the case that Q. This reasoning is not guaranteed to be valid. It confuses a sufficient condition for a necessary one.

Affirming the Consequent It is a mistake to conclude that the antecedent of a conditional is true just because the consequent is true. An argument of this form is not guaranteed to be valid: If P, then Q; Q is true; So, P is true too. This mistake involves mistaking a necessary condition for a sufficient one.

Red Herring Fallacy It is wrong to introduce something irrelevant in a discussion in order to divert attention away from an author's views or reasons. We saw that this can occur when, in the assessment of an author's reasoning about alternatives, a possibility that is unreasonable or highly unlikely is raised to show that an author is relying on a disjunction that is not exhaustive.

Appeal to Ignorance It is wrong to believe or decide that something is true (or false) just because it has not been proven otherwise. This can occur when reasoning about alternatives if one believes or judges that a disjunction is true or exhaustive just because one does not know of, or cannot think of, any other possibility. It also happens when one concludes that a claim is true on the grounds that one has not found a counterexample to it.

Equivocation It is a mistake in an argument when a word or phrase has to mean one thing for the argument to be valid and something else for the premises to be true. We also saw that it is a mistake when participants in a debate or discussion mean different things by some of the words they use. This is common in debates about pressing moral and political issues. It can be avoided or remedied by a careful use of the four-step method for defining words and claims, and by asking open-ended clarification questions.

False Conditional (Slippery Slope) It is a mistake to reason with a false or unjustified causal conditional. It is always, of course, a mistake for the premises of an argument not to be acceptable. A slippery slope argument is one in which it is argued that something will have unacceptable consequences. This kind of reasoning plays an important role in deciding what to do, since this invariably involves balancing the costs and benefits of competing proposals. But it is good reasoning only when the causal claim (the one that claims that the consequences will follow) is true, and only when it is true that those consequences are indeed bad.

Post Hoc It is a mistake to conclude that one event caused another solely on the grounds that the first preceded the other. While it is true that a cause always precedes its effect, this is not a sufficient condition for causation. It might just be a coincidence that the events happened in succession, or it might be that both events share a cause. As we saw

in Chapter 5, there are several methods for identifying the causes of an event.

False Analogy It is a mistake for an argument by analogy to include a false analogical premise. An analogical premise claims that the analogues are alike in all relevant respects. There are two ways to be mistaken in making an analogical claim. One might be mistaken about which respects are the relevant ones, or one might be mistaken about whether the analogues are alike in those respects. You should accuse an argument of committing a false analogy, only if you can identify a relevant difference. An argument by analogy using samples that has a false or unacceptable analogical claim as its premise commits the mistake called “hasty generalization.”

Unrepresentative Sample It is a mistake to draw a conclusion about a target group based on an unrepresentative sample. A sample is representative of a target just in case the sample and target are alike in all the relevant respects. This is simply a special case of the mistake of a false analogy. There are several things to consider in assessing whether a sample is representative.

- a. Is it too small? To be representative, a sample has to be large enough. How large it has to be depends on how homogeneous the target group is with respect to the relevant property. The more homogeneous it is, the smaller the sample size can be.
- b. Was it randomly selected? A sample is randomly selected just in case every member of the target group has the same chance of being in the sample group. Randomly selected sample of the right size are guaranteed to be representative.
- c. Was it self-selected? A sample is self-selected just in case the members of the sample are in the sample because they decided to be in it. A self-selected sample of the right size might be representative, but it is not very likely. It is best to treat self-selected samples as nonrepresentative.

Fallacy of Idealization It is a mistake to ignore or overlook the respects in which a model or map of some phenomenon involves idealization. A model involves idealization when relevant features of the phenomena are deliberately ignored. This can be fine, but only so long as one keeps tracks of the idealization.

APPENDIX B

PRACTICAL STRATEGIES

Critical thinking is reasonable and reflective thinking aimed at deciding what to believe or what to do. It is reflective in part because it requires us to think about our reasons as reasons, and so to ask whether they are acceptable and sufficient. It is reasonable in part because it requires us to look for reasons for our decisions. Throughout this book we have identified strategies and tricks that can help us be more reflective and reasonable in our thinking. Here are some of them.

GENERAL PURPOSE CRITICAL THINKING STRATEGIES

The Rule of Threes The rule of threes has such broad application that it is difficult to state simply. Here are some applications of it. Find three alternative courses of action when trying to decide what to do. Look for three objections to a view you are defending. Think about a problem from three different perspectives before trying to solve it. Look for three examples when trying to define a concept. Looking for three—or, even better, five!—will help you become more reflective in your thinking by forcing you to think “outside the box.” Usually, finding one or

two of the things you are looking for is relatively easy; trying to find more may force you to think harder, which is almost always good.

Don't Personalize Reasons Reasons and evidence do not belong to anyone; they are **universal**. And whether they are good has nothing to do with who accepts them; they are **objective**. To avoid personalizing reasons, replace the following:

- a. What evidence do you have?
- b. What are your reasons?
- c. Why do you believe that?

with the following impersonal ones:

- a'. What evidence is there?
- b'. What reasons are there to believe that?
- c'. Why should we believe that?

Think Twice; Decide Once To paraphrase the old carpenter's motto (measure twice and cut once), it is best to think twice and decide once. We know from psychological experiments that people are reluctant to change their minds. Once our opinions are set, it seems to take a lot of doing to revise them. For one thing, people tend to privilege evidence that confirms their already existing beliefs over evidence that conflicts with it. They assume that evidence that conflicts with what they already believe is probably not reliable. For another thing, people tend to prefer the evidence they have to evidence they would have to do something to get. To protect against these built-in obstacles to critical thinking, it is better to make sure that one has enough of the right kind of evidence before one makes a decision. It is better to think twice and decide once, than to have to go back and revise one's decisions.

Withhold Disagreement and Agreement We all know what it feels like to explain our opinions and views to people we know disagree with us. It puts additional pressure on us, pressure that is not really helpful if the goal is to get clear on the truth of the matter. For this reason, it is best to withhold disagreement until the speaker has had a full opportunity to explain, defend, and support her view. If she feels rushed into explaining it, she might leave out some crucial distinctions, and this might have the effect of weakening what might in fact be a fairly

plausible position. Or she might skip over some of the premises or fail to mention a response to an objection. Your goal as a critical thinker in a discussion is to help the other people make a case for their views that is as clear, as strong and as complete as possible.

Keep Emotional Distance Being a critical thinker means that our beliefs should be based on epistemic reasons, and not on emotional or pragmatic ones. Basing one's beliefs on emotions rather than on epistemic reasons is a mistake, since how a claim or proposal makes one feel emotionally is not a reliable guide to whether the claim is true or the proposal good. Emotions can also make it difficult to collect the evidence we need, or even from investigating further. Emotions get in the way when we identify too much with our own opinions and beliefs or with our own methods for collecting or evaluating evidence. If I become too emotionally attached to my beliefs and opinions, then I may react negatively when someone asks me for my reasons, or when they raise objections to my belief or when they state their own alternative beliefs. I might feel that they are criticizing me and not just my beliefs. The same is true if I am asked to defend my assessment of the evidence or my use of different methods for collecting evidence. If I come to identify too closely with these particular methods for assessing and collecting evidence, if I come to think of my value as a researcher as tied into their value, then I will react to criticisms of them as if they were criticisms of me and my judgment. This feeling of being under attack might make me feel defensive, and this can prevent me from thinking critically about the issue at hand. The same is true when I ask someone for his or her reasons. This sort of question is easily taken as aggressive or combative, even when the intention is simply to consider the issue from all sides as thoroughly as possible.

Trust, but (Be Prepared to) Verify Most critical thinking theorists agree that it would be asking too much to require that before the evidence from some source can be accepted we must first know that the source is reliable. Instead, they recommend the following: evidence from some source is acceptable; unless one has reason to think the source is not reliable. Trusting our sources is a default right, as it were. But we should not let ourselves get carried away. For we know that some apparent sources are not reliable at all, and others even ones that are reliable can still yield mistaken evidence. To borrow Ronald Reagan's remark about the proper attitude to take to enemy superpowers: trust, but be prepared to verify.

STRATEGIES FOR BEING REFLECTIVE ABOUT MEANING

Testing for Conceptual Independence It is good to know how to test for conceptual independence. In Chapter 1, we considered Robert Ennis' definition of critical thinking as reasonable, reflective thinking aimed at deciding what to believe or what to do. We also considered the standard philosophical definition of knowledge as justified, true belief. When an idea or concept is analyzed into several parts or elements, it is always a good idea to ask whether those parts or elements are conceptually independent of one another. To do this, simply ask yourself whether you can think of an example of something that has some of the elements but not others. For instance, we noticed that simple arithmetical calculations are a kind of thinking aimed at deciding what to believe but are not reflective because they do not require thinking about the method one uses. This example shows that reflective thinking and thinking that is aimed at deciding what to believe or do are conceptually independent. In one of the chapter's exercises, you discovered that in the case of knowledge, a belief's being true is conceptually independent of its being justified. Whenever a concept or problem has elements or parts, ask: how are those parts related to one another?

Look for Counter-examples Looking for counter-examples is an important step in constructing or evaluating a proposed definition or in assessing the truth of a conditional. A counter-example is a case that shows that the definition is either too broad (includes things that it should not) or too narrow (excludes things that it should include). The case could be an actual one, or it could simply be a fictional one. If you think that a proposed definition is too broad or too narrow, then you need to present a counter-example and argue that it shows that the proposed definition is mistaken. If, on the other hand, you respond to an alleged counter-example to your proposed definition, then you have to either show that it is not a genuine counter-example to the definition or else revise the definition to include/exclude examples of that kind.

Four-Step Definition Method A definition should state the meaning as clearly as possible and in as short a sentence as possible. This statement can usually take the form of a **slogan**. A definition should **expand** on that statement by filling in some of the detail that inevitably will get left out of a succinct statement. Among other things, the elaboration might say something about how the different elements in the statement are related one to another. This should take no more than a few sen-

tences. A definition should provide an **example** or two, depending on the complexity of what is being defined. The example could be from real life or it could be fictional, so long as it is clear and uncontroversial. Finally, a definition should identify some **contrasting** cases with which the thing being defined might easily be confused.

Ask Open-Ended Clarification Questions When discussing topics with other people, ask them open-ended questions, not questions that allow a “Yes” or “No” answer. This will reduce the risk that superficial agreement will mask interesting and deep differences. Instead of asking:

“Do you think that ...”

“Do you agree that ...”

Ask:

“Why do you think that ...”

“What do you mean by ...”

“What reasons are there for thinking that ...”

STRATEGIES FOR ANALYZING REASONS AND ARGUMENTS

Be Charitable When reconstructing someone’s reasons, it is best to try to turn it into a valid argument. This may require adding a missing premise. The reason to reconstruct arguments so that they are valid is that this focuses the discussion on whether the premises (including the ones that you added) are true, as opposed to the question of whether the premises are sufficient. As we know, adding a premise can make any argument valid. But this alone will not make an argument a good one: its premises must be true or acceptable as well.

The Assertion Test The premises in an argument are always assertions. But it is not always easy to tell in a passage or text just what the assertions are. To tell whether some proposition is asserted in a sentence of some text, ask whether the sentence as a whole could be true even if that proposition were false.

If the answer is Yes, then that proposition is not asserted.

If the answer is No, then that proposition is asserted.

Identifying Dependent Reasons It is important when analyzing a piece of reasoning to know whether the premises are working together (i.e., dependently) or whether they provide independent reason to accept the conclusion. There are two very useful strategies for deciding when the premises in an argument are dependent.

One is the **words test**: If the conclusion of an argument contains important words that occur only in one premise and important words that occur only in another premise, then those premises are probably dependent.

Another is the **false premise test**: If a premise would provide some reason to accept the conclusion even if another premise were *false*, then those premises are independent. So, to test whether premises are dependent, suppose that one is false, and ask whether the other one would still provide some support for the conclusion. If it would, then the premises are independent. If not, then they are dependent.

Testing for Logical Strength Test To test the logical strength of an argument, suppose that the premises were true. Then ask: how likely is that the conclusion would be true too? The higher the likelihood, the more logical support the premises provide.

INDEX

A

- Acceptability, questions about, 227
- Accurate representations, 206
- Addiction, meaning of, 66–68
- Advertising, 68–71, 124–126
- Affirming a disjunct, 155–156, 193
- Affirming the antecedent, 167
- Affirming the consequent, 170, 193
- Agreement, 184–185, 194, 233
- Alternatives, *see also* Disjunction(s)
 - acceptability of disjunctions, 152–154
 - criticizing reasoning about, 156–157
 - false alternative, 192
 - false disjunctions, 151–152
 - meaning of disjunctions, 149
 - questions about, 227
 - reasoning about, 149–158
- Analog(ies)
 - acceptability of, 200–204
 - analogical premise, 197, 207
 - analogues, 197
 - false analogy, 202, 217
 - models and maps, 213–215
 - reasoning by, 195–218
 - relevant similarity, 198–199
 - representational analogies, 205–216
 - samples, 207–213
 - validity of, 197–198
- Antecedent
 - affirming the, 167
 - and conditionals, 160–161
 - denying the, 169, 170, 193
- Appeals
 - to emotion, 32–33
 - to ignorance, 136, 153–154, 193, 224–225
 - to origins, 31
 - to relativism, 32
 - to tradition, 33–34
- Argument(s)
 - advantages of, 79
 - and critical thinking, 5, 75–80
 - defined, 77–78

Argument(s) (*cont'd*)

- diagramming of, 91
- from disagreement, 10–12
- features of, 97–98
- Logical Strength Test, 148
- and missing premises, 101–103
- subarguments, 94–96

Asserted noun clause, 48

Assertions

- Assertion Test, 44–48, 52, 72, 98–99, 149–150
- defined, 40–41
- and meaning, 41–48
- and premises and conclusions, 81–84
- and propositions, 41–44

Attitudes, 8–10

Autonomy, personal, 25–28

Averages, 132

B

Background checking, 128

Bad questions, 138, 144

Being Stuck in a Framework, 72

Belief(s)

- and arguments, 79
- clarifying of, 61–65
- and critical thinking, 13–14
- freedom of, 14
- justification of, 15–21
- and knowledge, 55–56
- and memory, 117
- and prejudice, 25–28

Bias

- and advertising, 125
- cultural bias, 129
- reporter bias, 128
- researcher bias, 141
- subject bias, 140
- and testimonial evidence, 122–123

Bi-conditionals, 162

Broad definitions, 52–53

C

Causal conditions, *see also*

Conditions

- acceptability of claims, 183–190
- concomitant variation, 186–187
- counterfactual conditionals, 181
- experimenting and simulating, 187–190
- false conditionals, 182–183, 194
- meaning of, 180–181
- necessary causal conditions, 184–185
- reasoning about, 179–190
- reasoning with, 181–183
- sufficient causal conditions, 185–186

Causal verbs, 179–180

Charged words, 139

Child kidnapping, 50

Circumstantial reasons, 91–93

Claims, causal, *see* Causal conditions

Clarification, asking for, 232–233

Cognitive illusions, 116

Common ground, 178

Community, 17

Comparison question, 201–202

Conceptual independence, 34

Conceptual relations, 56

Conclusion(s)

- conclusion indicators, 83–84
- defined, 77
- identification of, 80–86
- and the Words Test, 88

Concomitant variation, 186–187

Conditions, *see also* Causal conditions

- acceptability of claims, 173
- asserting a conditional, 46–47
- being explicit, 171–172
- bi-conditionals, 162
- and consequents, 160–161
- and defining words, 51–54
- invalid forms of reasoning, 169–171

- meaning of, 160–164
- necessary and sufficient conditions, 160–163, 186
- pure conditional reasoning, 168
- valid forms of reasoning, 165–168
- Confusing the value and the likelihood of a cost or benefit, 143–144
- Confusing undermining and overriding evidence, 143
- Conjunctions, 45–46
- Consequent
 - affirming the, 170, 193
 - and conditionals, 47, 160–161
 - denying the, 167
- Consistency, 134–135
- Contrast, 57
- Controls, 187–188
- Corporate bias, 128
- Costs and benefits
 - appeal to ignorance, 224–225
 - and decision making, 135, 188–189
 - likelihood and value, 136, 143–144
- Counter-example Method, 53
- Counter-example Strategy, 173–174, 194
- Counterfactual conditional, 181
- Criminal justice, 50
- Critical thinking
 - and arguments, 75–80
 - and belief, 13–14
 - critical thinking questions, 227–228
 - in decision making, 228–231
 - and disciplines, 220–227
 - in discussion, 232–234
 - and justification, 15–21
 - and knowledge, 6–24
 - nature and value of, 1–28
 - and personal autonomy, 25–28
 - and truth, 7–12
- Cultural bias, 129
- Culture, 17, 129
- D
- Day reconstruction method, 140
- Deciding what to do, *see* Decision making
- Decision making
 - and alternatives, 157, 231
 - comparing consequences, 188–189
 - costs and benefits, 135–136, 143–144, 188–189, 224–225
 - and critical thinking, 3
 - critical thinking in, 228–231
 - evaluating reasons, 100
 - measure twice, decide once, 144
 - and prejudice, 26–28
 - and reasons, 92–93, 230–231
 - think twice; decide once, 34–35
- Definition(s), *see also* Language; Word meaning
 - broad definitions, 52–53
 - constructing and evaluating, 49–58
 - contrast, 57
 - of critical thinking, 2–5
 - defining words, 51–56
 - definitional method. *see* SEEC method
 - examples, 56–57
 - false definition, 71
 - of key terms, 222–223
 - and meaning, 39–41
 - narrow definitions, 52–53
 - and necessary and sufficient conditions, 51–54
 - and perspectives, 40
 - reasoning with, 174–179
 - SEEC method, 50–51, 62–65, 72, 177
 - technical definitions, 66–67
- Degree, matter of, 110
- Denying a disjunct, 150
- Denying the antecedent, 169, 170, 193
- Denying the consequent, 167
- Difference, 185–186, 194
- Direct objections, 231

Disagreement, 10–12, 233

Disciplines

and critical thinking, 220–226

and definitions, 40

and frameworks, 59–61

history, 181

key concepts of, 220–223

legal theory, 203

modes of reasoning, 225–226

science, 120–121

sources of evidence, 223–225

Discussions, 71, 232–235

Disjunction(s), *see also* Alternatives

acceptability of, 152–154

affirming, 155–156, 193

asserting a disjunction, 46

criticizing reasoning about,
156–157

denying a disjunct, 150

exclusive disjunctions, 154–156

exhaustive disjunctions, 152–154

false disjunctions, 151–152

lucky disjunction, 153, 192

meaning of, 149

meaning of disjunctions, 149

symbolic form of, 151

Double blind study, 141

E

Emotion

appeal to emotion, 32–33

emotional distance, 234–235

emotional reasons, 16

Ends, deciding on, 19

Epistemic reasons, *see* Evidence

Equivocation, 176–177, 193

Evidence

acceptable evidence, 21–22

conflicting evidence, 22–24,
111–112

and disciplines, 223–225

and good reasons, 21–22

inappropriate evidence, 126

measurement, 131–142

and memory, 117–118

privileging available evidence, 33

reliability, 109–111

samples, 208–213

surveys, 136–141

testimonial evidence, 119–130

undermining and overriding

evidence, 22–23, 111–114, 143

witnesses, 119, 121–124, 127

Examples, 53, 56–57, 173–174, 194

Exhaustive disjunctions, 152–154

Experience, 121

Experimentation, 187–190, 194

Experts, 120–121, 125, 127

F

Fact checking, 128

Fallacies, *see* Mistakes

Fallacy of bad question, 138, 144

Fallacy of idealization, 214–215, 217

False alternatives, 152, 192

False analogy, 202, 217

False conditional, 182–183, 194

False premise test, 88–90, 105,
150, 166

Four-step method, *see* SEEC
method

Frameworks, 59–61, 72

Freedom, 14, 15

Fundamental concepts, 40

G

Goals, 19

Government policy, 60–61, 69–70

Government regulation, 68–71,
124–125

H

Happiness, 137–141

Hasty generalization, 213

Heritage, 17

History, 181

I

Idealization, 214–215, 217

If, then assertions, *see* Conditions

- Ignorance
 appeal to, 136, 153–154, 193,
 224–225
 relying on, 113, 153–154
- Illusions, 116
- Incommensurability, 135–136
- Incomplete information, 135
- J**
- Justification, 15–21, 55–56
- K**
- Key concepts, 220–223
- Key terms, 221
- Knowledge
 and belief, 13–15
 and critical thinking, 6–24
 definition of, 55–56
 and good reasons, 15–24
 and justification, 15–20
 knowledge chart exercise, 7
 and truth, 7–12
- L**
- Labels, 69–70
- Language, *see also* Word meaning
 and advertising, 68–71
 antecedent, 46
 and assertions and propositions,
 42–44
 causal verbs, 179–180
 charged words, 139
 conclusion indicators, 83–84
 conditionals, 46–47
 conjunctions, 45–46
 consequent, 47
 disjunction, 46
 noun clause, 47–48
 and reasoning, 41–42
 sentences, 42–44, 45–46
- Legal theory, 203
- Likelihood and value, 136,
 143–144
- Live possibilities, 157
- Loaded questions, 139
- Logic
 logical properties, 11
 Logical strength test, 98–99,
 105, 148
 logical support, 97–101
- Lost opportunities, 189
- Lucky disjunction, 192
- M**
- Mean, 132
- Meaning, 39–74, *see also* Word
 meaning
 in advertisements, 68–71
 assertions, 41–48
 beliefs and problems, 61–65
 and definitions, 39–41, 49–58
 frameworks, 59–61
 questions about, 227
 technical definitions, 66–67
- Means, 19
- Measurement
 acceptable reasons, 131–142
 consistency, 134–135
 evidence from, 131–134
 measurement fallacies, 143
 measure twice, decide once, 144
 precision, 135
 surveys, 136–141
- Median, 132
- Memory, 117–118
- Method
 of agreement, 184–185, 194
 concomitant variation, 194
 of difference, 185–186, 194
- Mistakes
 affirming a disjunct, 155–156, 193
 affirming the consequent,
 170, 193
 appealing to ignorance, 136,
 153–154, 193, 224–225
 appeal to emotion, 32–33
 appeal to origins, 31
 appeal to relativism, 32
 appeal to tradition, 33–34
 bad questions, 138, 144

Mistakes (*cont'd*)

- confusing meanings of “wrong,” 71
 - confusing the value and the likelihood of a cost or benefit, 143–144
 - confusing undermining and overriding evidence, 143
 - denying the antecedent, 169, 170, 193
 - equivocation, 176–177, 193
 - false alternatives, 192
 - false analogy, 202, 217
 - false conditional, 182–183, 194
 - false definition, 71
 - false dilemma, 151–152
 - false disjunction, 151–152
 - idealization, 214–215, 217
 - lucky disjunction, 153, 192
 - measurement fallacies, 143
 - personalizing reasons, 32
 - post hoc fallacy, 189–190, 194
 - privileging available evidence, 33
 - red herring fallacy, 157, 193
 - strawman mistake, 65, 72, 178, 193
 - talking past each other, 71
 - testimonial fallacy, 143
 - unrepresentative sample, 217
- Mode, 132
- Models and maps, 213–215
- Modus ponens, 166–167
- Modus tollens, 167, 168
- Morality, 177
- Multiple causes, 185
- N
- Narrow definitions, 52–53
- News reports, 126–129
- Nihilism, 8–10
- Noun clause, 47–48
- O
- Objections, 231
- Observation, 114–117
- Open-ended clarification questions, 63, 232–233
- Open-ended problems, 5, 41

- Opportunity cost, 189
 - Optical illusions, 116
 - Optimal conditions, 110
 - Origins, appeal to, 31
 - Overlooked conditions, 184–185
 - Overlooked effects, 185
 - Overriding evidence, 22–23, 111–114, 143
- P
- Patriot Act (U.S.), 69
- Perceptual observation, 115–117
- Personal autonomy, 25–28
- Personalizing reasons, 32
- Perspectives, 40, 59–61
- “Piling on” effect, 91–92
- Possibilities, 149–158
- Post hoc fallacy, 189–190, 194
- Pragmatic reasons, 16
- Precision, 135
- Prejudice, 25–28
- Premises
- circumstantial reasons, 91–93
 - defined, 77
 - dependent and independent reasons, 87–93
 - false premise test, 88–90, 105, 150, 166
 - identification of, 80–86
 - missing premises, 101–103
 - premise indicators, 82, 84
 - subarguments, 94–96
 - words test, 87–88
- Privileging available evidence, 33
- Probability, 136
- Problems, 5, 41, 61–65
- Proposals, 64–65
- Propositions, 41–44
- Psychology, 118
- Pure conditional reasoning, 168
- Q
- Quantifiers, 207–208
- Questions
- bad questions, 138, 144
 - comparison questions, 201–202

- in discussions, 232–233
 - kinds of, 227–228
 - loaded questions, 139
 - open-ended clarification questions, 63, 232–233
 - relevancy questions, 201
- R**
- Realism, 8–10
 - Reasonable thinking, 4
 - Reasons, *see also* Premises
 - and critical thinking, 15–21
 - emotional reasons, 16
 - epistemic reasons, 17–20
 - good reasons, 21–25
 - and measurement, 131–142
 - memory, 117–118
 - observation, 114–117
 - personalizing reasons, 32
 - pragmatic reasons, 16
 - reliable evidence, 109–114
 - sufficient reasons, 75–108
 - and testimony, 119–130
 - Red herring fallacy, 157, 193
 - Referents, 51
 - Reflective thinking, 4–5
 - Relativism, 8–12, 32
 - Relevancy questions, 201
 - Relevant property, 197, 199–200
 - Reliability, 109–114, 117–118, 133, 134–135
 - Religious studies, 211–212
 - Reporters, 126–127, 128
 - Representations, accurate, 206
 - Research
 - double blind study, 141
 - experimentation, 187–190, 194
 - fact checking, 128
 - measurement, 131–135
 - researcher bias, 141
 - samples, 208–213
 - surveys, 136–141
 - Responsibility, 15
 - Rule of threes, 60, 228
- S**
- Samples, 208–213
 - random samples, 209–210
 - reasoning with, 207
 - sample group, 207
 - sample size, 208–209
 - self-selected samples, 210–213
 - unrepresentative sample, 217
 - Scales, 132–133, 138–139
 - Science, 120–121
 - SEEC method
 - for clarification, 62–65
 - definitional method, 177
 - described, 50–51
 - Four-Step Method, the, 72
 - and proposals, 64–65
 - Sensory observation, 115
 - Sentences, *see* Language
 - Simulation, 194
 - Simulations, 189
 - Slanted words, 139
 - Slippery slope fallacy, 182–183, 194
 - Slogans, 51–54, *see also* Word meaning
 - Social relativist, 9
 - Sources, 128, 223–225
 - Standardized scales, 132–133, 138–139
 - Standards, 174–179
 - Strawman mistake, 65, 72, 178, 193
 - Subarguments, 94–96
 - Subject bias, 140
 - Subjective relativist, 9
 - Sufficiency
 - and arguments, 75–80
 - dependent and independent reasons, 87–93
 - logical support, 97–101
 - missing premises, 101–103
 - premises and conclusions, 80–86
 - questions about, 227
 - subarguments, 94–96
 - sufficient evidence, 21–22
 - sufficient reasons, 75–108

Sufficient conditionals, *see*
 Causal conditions;
 Conditions
 Surveys, 136–141

T

Taking agreement for
 granted, 233
 Talking past each other, 71
 Target group, 207
 Taste, matters of, 120
 Tax relief, 69–70
 Technical definitions, 66–67
 Testimonial evidence, *see also*
 Evidence
 acceptable reasons, 119–130
 advertising, 124–126
 appropriate testimony,
 120, 126
 competent testimony, 121
 defined, 119
 news reports, 126–129
 and surveys, 139
 testimonial fallacy, 143
 unacceptable testimony,
 122–123
 unbiased testimony, 122–123
 unjustified criticism of, 129
 witnesses, 117–118, 119,
 121–124, 127
 Therefore Test, the, 83–84, 85
 Think twice; decide once, 34–35
 Threes, rule of, 60, 228
 Topic relativity, 110
 Tradition, appeal to, 33–34
 Training, 121
 Trust, but (be prepared to) verify,
 144
 Truth
 defining of, 7–12
 and knowledge, 55–56
 nihilism, 9
 questions about, 227
 realism, 8
 relativism, 8–9

U

Unacceptable testimony, 122–123
 Unbiased testimony, 122–123
 Undermining evidence, 23–24,
 111–114
 Unjustified criticism of testimony,
 129
 Utilitarians, 137

V

Validity
 definition of, 222–223
 and missing premises, 101–103
 and truth, 99
 and uncertainty, 203
 valid, defined, 11
 Values, 177–178, 228

W

Witnesses, *see also* Evidence;
 Testimonial evidence
 and bias, 122–123
 and competent testimony, 121
 and memory, 117–118
 and news reports, 127
 in trials, 119
 undermining of, 123–124
 Word, referents of, 51
 Word meaning, *see also*
 Definition(s); Language;
 Meaning
 and advertisements, 68–71
 and conditionals, 174–179
 disagreements about, 66–68
 equivocation, 176–177
 narrow and broad definitions,
 52–53
 necessary and sufficient
 conditions, 51–52
 slanted words, 139
 slogans, 51–55
 of “wrong,” 71
 Words Test, 87–88, 105
 Writing arguments, 77–78
 Wrong, meaning of, 54, 71