Epistemology

A Contemporary Introduction to the Theory of Knowledge

Third Edition

Robert Audi

Routledge
Taylor & Francis Group
NEW YORK AND LONDON

Routledge Contemporary Introductions to Philosophy
Series editor: *Paul K Moser*, Loyola University of Chicago

This innovative, well-structured series is for students who have already done an introductory course in philosophy. Each book introduces a core general subject in contemporary philosophy and offers students an accessible but substantial transition from introductory to higher-level college work in that subject. The series is accessible to non-specialists and each book clearly motivates and expounds the problems and positions introduced. An orientating chapter briefly introduces its topic and reminds readers of any crucial material they need to have retained from a typical introductory course. Considerable attention is given to explaining the central philosophical problems of a subject and the main competing solutions and arguments for those solutions. The primary aim is to educate students in the main problems, positions, and arguments of contemporary philosophy rather than to convince students of a single position.

Epistemology
Third edition
Robert Audi

Philosophy of Mind
Third edition
John Heil

Philosophy of Perception
William Fish

Metaphysics
Third edition
Michael J. Loux

Philosophy of Mathematics
Second edition
James Robert Brown

Classical Philosophy
Christopher Shields

Classical Modern Philosophy
Jeffrey Tlumak

Ethics
Harry Gensler

Philosophy of Science
Second edition
Alex Rosenberg

Philosophy of Biology
Alex Rosenberg and Daniel W. McShea

Philosophy of Language
Second edition
Willam G. Lycan

Philosophy of Religion
Keith E. Yandell

Philosophy of Art
Noël Carroll

Philosophy of Psychology
José Bermudez

Social and Political Philosophy
John Christman

Continental Philosophy
Andrew Cutrofello

Free Will (forthcoming)
Michael McKenna

Ethics (forthcoming)
Second edition
Harry Gensler

Feminist Philosophy (forthcoming)
Lorraine Code

Philosophy of Literature (forthcoming)
John Gibson

First edition published 1998, second edition published 2003
by Routledge
This edition first published 2011
by Routledge
270 Madison Ave, New York, NY 10016

Simultaneously published in the UK
by Routledge
2 Park Square, Milton Park, Abingdon, Oxon OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

This edition published in the Taylor & Francis e-Library, 2010.

To purchase your own copy of this or any of Taylor & Francis or Routledge's collection of thousands of eBooks please go to www.eBookstore.tandf.co.uk.

Library of Congress Cataloging in Publication Data
Audi, Robert
Epistemology : a contemporary introduction to the theory of knowledge / Robert Audi. — 3rd ed.
p. cm. — (Routledge contemporary introductions to philosophy)
Includes bibliographical references and index.
[etc.]
1. Knowledge, Theory of. I. Title.
BD161.A783 2010
121—dc22
2010005139

ISBN 13: 978-0-415-87922-4 (hbk)
ISBN 13: 978-0-415-87923-1 (pbk)

To Malou

Contents

Preface to the first edition

This book is a wide-ranging introduction to epistemology, conceived as the theory of knowledge and justification. It presupposes no special background in philosophy and is meant to be fully understandable to any generally educated, careful reader, but for students it is most appropriately studied after completing at least one more general course in philosophy.

The main focus is the body of concepts, theories, and problems central in understanding knowledge and justification. Historically, justification—sometimes under such names as 'reason to believe', 'evidence', and 'warrant'—has been as important in epistemology as knowledge itself. This is surely so at present. In many parts of the book, justification and knowledge are discussed separately; but they are also interconnected at many points. The book is not historically organized, but it does discuss selected major positions in the history of philosophy, particularly some of those that have greatly influenced human thought. Moreover, even where major philosophers are not mentioned, I try to take their views into account. One of my primary aims is to facilitate the reading of those philosophers, especially their epistemological writings. It would take a very long book to discuss representative contemporary epistemologists or, in any detail, even a few historically important epistemologies, but a shorter one can provide many of the tools needed to understand them. Providing such tools is one of my main purposes.

The use of this book in the study of philosophy is not limited to courses or investigations in epistemology. Epistemological problems and theories are often interconnected with problems and theories in the philosophy of mind; nor are these two fields of philosophy easily separated (a point that may hold, if to a lesser extent, for any two central philosophical fields). There is, then, much discussion of the topics in the philosophy of mind that are crucial for epistemology, for instance the phenomenology of perception, the nature of belief, the role of imagery in memory and introspection, the variety of mental properties figuring in self-knowledge, the nature of inference, and the structure of a person's system of beliefs.

Parts of the book might serve as collateral reading not only in pursuing the philosophy of mind but also in the study of a number of philosophers often discussed in philosophy courses, especially Plato, Aristotle, Aquinas,

Descartes, Leibniz, Locke, Berkeley, Hume, Kant, and Mill. The book might facilitate the study of moral philosophy, such as Kantian and utilitarian ethics, both discussed in some detail in Chapter 9; and it bears directly on topics in the epistemology of religion, some of which are also discussed in Chapter 9.

The writing is intended to be as simple and concrete as possible for a philosophically serious introduction that does not seek simplicity at the cost of falsehood. The territory surveyed, however, is extensive and rich. This means that the book cannot be traversed quickly without missing landmarks or failing to get a view of the larger segments and their place in the whole. Any one chapter can perhaps be read at a sitting, but experience has shown that even the shortest chapter covers too many concepts and positions for most readers to assimilate in a single reading and far more than most instructors can cover in any detail in a single session.

To aid concentration on the main points, and to keep the book from becoming more complicated, notes are limited, though parenthetical references are given in some places and there is also a short selected bibliography with thumbnail annotations. By and large, the notes are not needed for full comprehension and are intended mainly for professional philosophers and serious students. There are also some subsections that most readers can probably scan, or even skip, without significant loss in comprehending the *main* points of the relevant chapter. Technical terms are explained briefly when introduced and are avoided when they can be. Most of the major terms central in epistemology are defined or explicated, and boldfaced numbers in the index indicate main definitional passages. But some are indispensable: they are not mere words, but tools; and some of these terms express concepts valuable outside epistemology and even outside philosophy. The index, by its boldfaced page references to definitions, obviates a glossary.

It should also be stressed that this book is mainly concerned to introduce the *field* of epistemology rather than the *literature* of epistemology—an important but less basic task. It will, however, help non-professional readers prepare for a critical study of that literature, contemporary as well as classical. For that reason, too, some special vocabulary is introduced and a number of the notes refer to contemporary works.

The sequence of topics is designed to introduce the field in a natural progression: from the genesis of justification and knowledge (Part One), to their development and structure (Part Two), and thence to questions about what they are and how far they extend (Part Three). Even apart from its place in this ordering, each chapter addresses a major epistemological topic, and any subset of the chapters can be studied in any order provided some appropriate effort is made to supply the (generally few) essential points for which a later chapter depends on an earlier one.

For the most part this book *does* epistemology rather than talk about it or, especially, about its literature. In keeping with that focus, the ordering of chapters is intended to encourage *understanding* epistemology before

discussing it in large-scale terms, for instance before considering what sort of epistemological theory, say normativist or naturalistic, best accounts for knowledge. My strategy is, in part, to discuss myriad cases of justification and knowledge before approaching analyses of what they are, or the skeptical case against our having them.

In one way, this approach differs markedly from that of many epistemological books. I leave the assessment of skepticism for the last chapter; early passages indicate that skeptical problems must be faced and, in some cases, how they are connected with the subject at hand or are otherwise important. Unlike some philosophers, I do not think extensive discussion of skepticism is the best way to motivate the study of epistemology. Granted, historically skepticism has been a major motivating force; but it is not the only one, and epistemological concepts hold independent interest. Moreover, in assessing skepticism I use many concepts and points developed in earlier chapters; to treat it early in the book, I would have to delay assessing it.

There is also a certain risk in posing skeptical problems at or near the outset: non-professional readers may tend to be distracted, even in discussing conceptual questions concerning, say, what knowledge *is*, by a desire to deal with skeptical arguments purporting to show that there is none. There may be no best or wholly neutral way to treat skepticism, but I believe my approach to it can be adapted to varying degrees of skeptical inclination. An instructor who prefers to begin with skepticism can do so by taking care to explain some of the ideas introduced earlier in the book. The first few sections of Chapter 10 (Chapter 13 in the third edition), largely meant to introduce and motivate skepticism, presuppose far less of the earlier chapters than the later, evaluative discussion; and most of the chapter is understandable on the basis of Part One, which is probably easier reading than Part Two.

My exposition of problems and positions is meant to be as nearly unbiased as I can make it, and where controversial interpretations are unavoidable I try to present them tentatively. In many places, however, I offer my own view. Given the scope of the book, I cannot provide a highly detailed explanation of each major position discussed, or argue at length for my own views. I make no pretense of treating anything conclusively. But in some cases—as with skepticism—I do not want to leave the reader wondering where I stand, or perhaps doubting that there is any solution to the problem at hand. I thus propose some tentative positions for critical discussion.

Acknowledgments to the first edition

This book has profited from my reading of many articles and books by contemporary philosophers, and from many discussions I have had with them and, of course, with my students. I cannot mention all of these philosophers, and I am sure that my debt to those I will name—as well as to some I do not, such as some whose journal papers I have read but have not picked up again, and some I have heard at conferences—is incalculable. Over many years, I have benefited greatly from discussions with William Alston, as well as from reading his works; and I thank him for detailed critical comments on parts of the manuscript. Reading of books or articles (or both) by Roderick Chisholm, Richard Foley, Paul Moser, Alvin Plantinga, Walter Sinnott-Armstrong, and Ernest Sosa, and a number of discussions with them, have also substantially helped me over many years. My colleagues at the University of Nebraska, especially Albert Casullo, and several of my students have also helped me at many points. I have learned greatly from the participants in the National Endowment for the Humanities seminars and institutes I have directed. I also benefited much from the papers given to the seminars or institutes by (among others) Laurence BonJour, Fred Dretske, Alvin Goldman, Gilbert Harman, Keith Lehrer, Ruth Marcus, and John Perry, with all of whom I have been fruitfully discussing epistemological topics on one occasion or another for many years.

In relation to some of the main problems treated in the book, I have learned immensely from many other philosophers, including Frederick Adams, Robert Almeder, David Armstrong, John A. Barker, Richard Brandt, Panayot Butchvarov, Carol Caraway, the late Hector-Neri Castañeda, Wayne Davis, Michael DePaul, Susan Feagin, Richard Feldman, Roderick Firth, Richard Fumerton, Carl Ginet, Alan Goldman, Risto Hilpinen, Jaegwon Kim, John King-Farlow, Peter Klein, Hilary Kornblith, Christopher Kulp, Jonathan Kvanvig, Brian McLaughlin, George S. Pappas, John Pollock, Lawrence Powers, W.V. Quine, William Rowe, Bruce Russell, Frederick Schmitt, Thomas Senor, Robert Shope, Donna Summerfield, Marshall Swain, William Throop, Raimo Tuomela, James Van Cleve, Thomas Vinci, Jonathan Vogel, and Nicholas Wolterstorff. In most cases I have not only read some

epistemological work of theirs, but discussed one or another epistemological problem with them in detail.

Other philosophers whose comments or works have helped me with some part of the book include Anthony Brueckner, Stewart Cohen, Earl Conee, Dan Crawford, Jonathan Dancy, Timothy Day, Robert Fogelin, Elizabeth Fricker, Bernard Gert, Heather Gert, David Henderson, Terence Horgan, Dale Jacquette, Eric Kraemer, Noah Lemos, Kevin Possin, Dana Radcliffe, Nicholas Rescher, Stefan Sencerz, James Taylor, Paul Tidman, Mark Timmons, William Tolhurst, Mark Webb, Douglas Weber, Ümit Yalçin, and Patrick Yarnell.

I owe special thanks to the philosophers who generously commented in detail on all or most of some version of the manuscript: John Greco, Louis Pojman, and Matthias Steup. Their numerous remarks led to many improvements. Detailed helpful comments were also provided by readers for the Press, including Nicholas Everett, Frank Jackson, and Noah Lemos. All of the philosophers who commented on an earlier draft not only helped me eliminate errors, but also gave me constructive suggestions and critical remarks that evoked both clarification and other improvements. I am also grateful for permission to reuse much material that appears here in revised form from my *Belief, Justification, and Knowledge* (Wadsworth Publishing Co., 1988) and I thank the editor of *American Philosophical Quarterly* for allowing me to use material from 'The Place of Testimony in the Fabric of Knowledge and Justification' (vol. 34, 1997). For advice and help at several stages I thank Paul Moser, Editor of the series in which this book appears, and Adrian Driscoll and the staff at Routledge in London, including Moira Taylor and Sarah Hall, and Dennis Hodgson.

Robert Audi
February, 1997

Preface to the second edition

This preface will presuppose the Preface to the first edition and can therefore be brief. Many improvements have been made in this edition, but they do not make the previous preface inapplicable, and reading it should help anyone considering a study of even part of the book.

My main concern in revising has been to produce a book that is *both* philosophically stronger and easier to read. Doing this has required adding new substantive material, making minor changes throughout, adding or extending many examples, making various refinements and corrections, and bringing in new references, notes, and bibliography.

Instructors who have used the volume in their teaching will find that the content and organization are highly similar and that a transition from the first edition to this one is easy. Students and people reading for general interest should find the book easier to understand. The emphasis is still on enhancing comprehension of the *field* of epistemology—its concepts, problems, and methods—rather than on presenting its literature. But, perhaps even more than in the first edition, the book is generally in close contact with both classical and contemporary literature. In this edition there are also many more references to pertinent books and papers, particularly those published in recent years.

This edition includes more extensive discussion of *virtue epistemology* and *social epistemology*, with *feminist epistemology* figuring significantly (though not exclusively) in relation to social epistemology. The connection of epistemology with philosophy of mind and language also receives more emphasis in this edition. So does *contextualism* and the related theory of "relevant alternatives."

I am happy to say that Routledge has published a fine and wide-ranging new collection of readings to accompany this book: Michael Huemer's *Epistemology: Contemporary Readings* (2002). Huemer has chosen classical and contemporary book sections and papers that go well with every chapter in the present book; his larger sections match mine; and he offers helpful introductions to each section and study questions on each chapter.

This edition of my book is certainly self-contained, but its integration with Huemer's supporting collection (for which I have done a long narrative introduction to help both instructors and students) is close, and the two together provide enough substance and diversity to facilitate numerous different kinds of epistemology courses.

Acknowledgments to the second edition

Fortunately, I have continued to benefit from reading of or discussions with most of the people acknowledged above, in the first edition. But since that writing I have come to know the work of many other writers in the field and profited from teaching many more students. I am bound to omit some people I should thank, but I want to acknowledge here a number of people not mentioned above: John Broome, Tyler Burge, David Chalmers, Roger Crisp, Mario DeCaro, Keith DeRose, Rosaria Egidi, Guido Frongia, Douglas Geivett, Joshua Gert, Peter Graham, D.W. Hamlyn, Brad Hooker, Christopher Hookway, Michael Huemer, Jonathan Jacobs, Ralph Kennedy, Simo Knuutila, Matthias Lutz-Bachmann, Hugh McCann, John McDowell, Tito Megri, Cyrille Michon, Nicholas Nathan, Ilkka Niiniluoto, Tom O' Neil, Derek Parfit, James Pryor, Jlenia Quartarone, Joseph Raz, John Searle, John Skorupski, David Sosa, William Talbot, Wilhelm Vossenkuhl, Fritz Warfield and Paul Weirich. In addition, I have continued to benefit from regular exchanges of ideas or papers (usually both) with many philosophers, including William Alston, Laurence BonJour, Panayot Butchvarov, Elizabeth Fricker, Alvin Goldman, John Greco, Gilbert Harman, Jaegwon Kim, Christopher Kulp, Jonathan Kvanvig, Bruce Russell, Walter Sinnott-Armstrong, Ernest Sosa, Matthias Steup, Eleonore Stump, Richard Swinburne, Raimo Tuomela, Nicholas Wolterstorff, and Linda Zagzebski.

Readers of earlier versions of this edition deserve special thanks, not only the anonymous readers for the Press but also Michael Pace, Bruce Russell, Mark Owen Webb, and especially Claudio de Almeida, who provided numerous expert comments and criticism (more, indeed, than I could fully take into account in the available time and space). Special thanks are also due the Editors at Routledge, particularly Tony Bruce, Simon Bailey, and Siobhan Pattinson, whose ideas and support have been immensely helpful.

Preface to the third edition

This edition reflects some of the benefits of nearly a decade of teaching and writing about epistemology since the second edition went to press. There are revisions and improvements throughout. Some revisions are responsive to new developments in epistemological literature; others respond to comments by professional colleagues, including various instructors who have used the book. Student responses have also been taken into account.

The book is structurally much as in the second edition, and the previous prefaces largely apply to it. As before, instructors should read prefatory material and the introduction. Long chapters have been divided, but the chapters remain cumulative in content. Most of them, however, can be read independently of the others or by simply looking into some earlier chapter at certain points. The index may also help readers, and its boldface numerals indicate places where the term in question is defined.

Those who have used the second edition will find no difficulty adjusting their teaching or discussions to this one. Chapters that have been divided still cover the same issues, though with new material included and revisions of much that is included. The fit with Michael Huemer's large collection of major papers (cited in the bibliography) is equally good.

Some topics treated in this edition are not addressed in the second. These include the nature of intuitions, the skeptical challenge of rational disagreement, and the value problem: the range of questions concerning why knowledge and justified true belief have value beyond that of merely true belief. Other topics receive considerably more exploration than in the second edition, especially contextualism, perception (including perceptual content), self-evidence and the a priori, memorial justification, inferential versus direct knowledge, inference to the best explanation, scientific hypotheses, testimony and trust, understanding, and virtue epistemology.

Acknowledgments to the third edition

I have been fortunate in being able to exchange ideas with most of the philosophers named in the previous two sets of acknowledgments, and this book has continued to benefit from those discussions. I look back on them with much gratitude. I should particularly mention continuing conversations with the late William P. Alston and with Laurence BonJour, Mario DeCaro, Elizabeth Fricker, Peter Graham, John Greco, Ralph Kennedy, Peter Klein, Christopher Kulp, Jonathan Kvanvig, Jennifer Lackey, Bruce Russell, Walter Sinnott-Armstrong, Ernest Sosa, and Thomas Vinci.

On one or another topic in the book I have greatly benefited from epistemological discussions with my Notre Dame colleagues Marian David, Michael DePaul, Peter van Inwagen, Alvin Plantinga, Leopold Stubenberg, and Fritz Warfield. Interactions with Andrew Bailey, Michael Bergmann, E.J. Coffman, Roger Crisp, Claudio De Almeida, Keith DeRose, Fred Dretske, Richard Feldman, Richard Fumerton, Alvin Goldman, Christopher Green, Stephen Grimm, Michael Huemer, Thomas Kelly, Matthew Kennedy, Hilary Kornblith, Markus Lamenranta, Duncan Pritchard, Baron Reed, David Sosa, Jeff Speaks, Matthias Steup, Raimo Tuomela, Jonathan Vogel, Michael Zimmerman, and, especially, Sanford Goldberg and Timothy Williamson have also been of great help to me.

Detailed comments on the penultimate version were provided by Scott Hagaman and (on a few parts) by Matthew Lee and Ezra Cook, and these were of great help. Hagaman also made numerous helpful comments and critical remarks in seminar discussions at Notre Dame. The numerous questions and responses by seminar students at the Universities of Notre Dame and Nebraska have been valuable at many points.

The editorial advice of Andrew Beck, Andrew R. Davidson, Michael Andrews, and, since the book's first edition, Tony Bruce have helped me in several important matters. They and their staff have been essential in both design and production.

Introduction

A sketch of the sources and nature of belief, justification, and knowledge

Before me is a grassy green field. A line of trees marks its far edge, which is punctuated by a spruce on its left side and a maple on its right. Birds are singing. A warm breeze brings the smell of roses from a nearby trellis. I reach for a glass of iced tea, still cold to the touch and flavored by fresh mint. I am alert, the air is clear, the scene is quiet. My perceptions are quite distinct.

It is altogether natural to think that from perceptions like these, we come to know a great deal—enough to guide us through much of daily life. But we sometimes make mistakes about what we perceive, just as we sometimes misremember what we have done, or infer false conclusions from what we believe. We may then think we know something when in fact we do not, as when we make errors through inattention or are deceived by vivid dreams. And is it not possible that we are mistaken more often than we think?

Perception, belief, and justification

Philosophers have thought a great deal about these matters, especially about the nature of perceiving and about what we can know—or may mistakenly think we know—through perception or through other sources of knowledge, such as memory as a storehouse of what we have learned in the past, consciousness as revealing our inner lives, reflection as a way to acquire knowledge of abstract matters, and testimony as providing knowledge originally acquired by others. In approaching these topics in epistemology—the theory of knowledge and justification—it is appropriate to begin with perception. In my opening description, what I detailed was what I *perceived*: what I saw, heard, smelled, felt, and tasted. In describing my experience, I also expressed some of what I *believed*: that there was a green field before me, that there were bird songs, that there was a smell of roses, that my glass felt cold, and that the tea tasted of mint.

It seems altogether natural to believe these things given my experience, and I think I *justifiedly* believed them. I believed them, not in the way I would if I accepted the result of wishful thinking or of merely guessing, but with justification. By that I mean above all that the beliefs I refer to were *justified*.

This a good thing; justified beliefs are of a kind it is desirable and reasonable to hold.

Justification as process, as status, and as property

Being justified, in the sense illustrated by my beliefs about what is clearly before me, need not be the result of a process. Being justified is not, for instance, like being purified, which requires a process of purification. My beliefs about what is before me are not justified because they have been through *a process of being justified*, as when we defend a controversial belief by giving reasons for it. They have not; the question whether they are justified has not even come up. No one has challenged them or even asked why I hold them. They are justified—in the sense that they have the *property of being justified—justifiedness*—because there is something about them in virtue of which they are natural and appropriate for me as a normal rational person.

We can see what justifiedness is by starting with a contrast. Unlike believing something that one might arrive at through a wild guess in charades, our justified perceptual beliefs are justified for us simply through their arising in the normal way from our clear perceptions. Roughly, they are justified in the sense that they are quite in order from the point of view of the standards for what we may reasonably believe. That, in turn, is roughly what we may believe without being subject to certain kinds of criticism, say as intellectually lax, as sloppy, as overhasty, or the like. Justified beliefs are also a kind that we tend to expect to be true. Imagine someone's saying 'His belief is justified, but I don't expect it to turn out to be true'. Without special explanation, this would be to take away with one hand something given by the other.

In saying that I justifiedly believe there is a green field before me, I am implying something else, something quite different, though it sounds very similar, namely that I am *justified in believing* there is a green field before me. To see the difference, notice that we can be justified *in* believing something—roughly in the sense that we *have* a justification for believing it—without believing it at all, quite as we can be justified in doing something, such as criticizing a person who has failed us, without doing it. Similarly, I might be justified in believing that I can do a certain difficult task, yet fail to believe this until someone helps me overcome my hesitation. I may then see that I *should* have believed it.

Being justified in believing something is having justification *for* believing it. This, in turn, is roughly a matter of having ground for believing it (and we also speak of having *a* ground or *a* justification or *a* reason). Just as we can have reason to do things we do not do, we can have reason to believe things we do not believe. You can have reason to go to the library and forget to, and I can have reason to believe someone is making excuses for me but—because I have no inkling that I need any—fail to believe this. Our justification for believing is basic raw material for actual justified belief; and justified belief is commonly good raw material for knowledge.

The two justificational notions are intimately related: if one justifiedly believes something, one is also justified *in* believing it, hence *has* justification for believing it. But the converse does not hold: not everything we are justified *in* believing is something we *do* believe. When I look at a lawn, I am justified *in* believing it has more than ten blades of grass per square foot, but I would not normally have any belief about the number of blades per square foot. We have more justificational raw material than we need or use. We do not believe anywhere near the number of things that we have justification to believe. This holds not just in trivial matters but also in, for instance, mathematics.

There are many things we are justified in believing which we do not actually believe, such as the proposition that normal people do not drink 100 liters of water a day. Let us call the first kind of justification—justifiedly believing—*belief justification*, as it belongs to actual beliefs. It is also called *doxastic* justification, from the Greek *doxa*, translatable as 'belief'. Call the second kind—being justified *in* believing—*situational justification*, since it is based on the informational situation one is in. It is a *status* one has in virtue of that situation. This situation includes not just what one perceives, but also one's background beliefs and knowledge, such as the belief that people drink at most a few liters of water a day. Situational justification is also called *propositional justification,* since the proposition in question is justified *for* the person whose situation provides justification for believing it, and the person *has* justification for it.

In any ordinary situation in waking life, we have both a lot of general information stored in memory and much specific information presented in our perceptions. We do not need all this information, and our situational justification for believing something is often unaccompanied by our actually believing that it is so. We have situational justification for vastly more justified beliefs than we actually have. Here nature is very generous. We are built to gain from a mere glance enough information to ground vastly more beliefs than we normally form or rely on.

Without situational justification, such as the kind that comes from seeing a green field, there would be no belief justification. I would not, for instance, justifiedly believe that there is a green field before me. We cannot have a justified belief without being in a *position* to have it. Without situational justification, we are not in such a position. Without belief justification, on the other hand (i.e., doxastic justification), we would have no beliefs of a kind we want and need, those with a positive status—being justified—that makes them appropriate for us as rational creatures and warrants us in expecting them to be true. Belief justification, then, is more than the situational kind it presupposes.

Belief justification occurs when there is a certain kind of connection between what yields situational justification and the justified belief that benefits from it. Belief justification occurs when a belief is *grounded in*, and thus in a way supported by (or *based on*), something that gives one situational

justification for that belief, such as seeing a field of green. Seeing is of course perceiving; and perceiving is a basic source of knowledge—perhaps our most elemental source, at least in childhood. This is largely why perception is so large a topic in epistemology.

Knowledge and justification

Knowledge would not be possible without belief justification—or a kind of grounding significantly like it. If I did not have the kind of justified belief I do—if, for instance, I were wearing dark sunglasses and could not tell the difference between a green field and a smoothly ploughed one that is really an earthen brown—then on the basis of what I now see I would not know that there is a green field before me.

To see how knowledge fits into the picture so far sketched, consider two points. First, justified belief is important for knowledge because at least the typical things we know we *also* justifiedly believe on the same basis that grounds our knowing them. If I know someone is making excuses for me, say by the way she explains my lateness, I do not just believe this but justifiedly believe it. Second, much of what we justifiedly believe we also know. Surely I could have maintained, regarding each of the things I have said I justifiedly believed through perception, that I also knew it. And do I not know these things—say that there is a lawn before me and a car on the road beyond it—on the *same* basis on which I justifiedly believe them, for instance on the basis of what I see and hear? This is very plausible.

As closely associated as knowledge and justified belief are, there is a major difference. If I know that something is so, then it is *true*, whereas I can justifiedly believe something false. If a normally reliable friend tricked me into believing something false, say that he lost my car keys, I could still justifiedly believe he lost them. We must not assume, then, that everything we learn about justified belief applies to knowledge. We should look at both concepts independently.

I said that I *saw* the green field and that my belief that there was a green field before me arose from my seeing it. If the belief arose, under normal conditions, from my seeing the field (so that I believed it is there simply because I saw it there), then the belief was true, justified, and constituted knowledge. Again, however, we can alter the example to bring out how knowledge and justification may diverge: the belief might remain justified even if, unbeknownst to me, the grass had been burned up since I last saw it, and there was now a perfect artificial replica of it spread out in grassy-looking strips of cloth that hide the charred ground. Then, although I might think I know the green field is there, I would only falsely believe I know this. Such a bizarre happening is, to be sure, improbable. Still, a justified but false belief *could* arise in this way.

Memory, introspection, and self-consciousness

As I look at the field before me, I *remember* carefully cutting a poison ivy vine from the trunk of the spruce. Surely, my memory belief that I cut off this vine is justified. I think I also know that I did this. But here I confess to being less confident than I am of the justification of my perceptual belief, held in the radiant sunlight, that there is (now) a green field before me.

As our memories become less vivid, we tend to be correspondingly less sure that our beliefs apparently based on them are justified. Still, I distinctly recall cutting the vine. The stem was furry; it was bonded to the tree trunk; the cutting was difficult and slightly wounded the tree. By contrast, I have no belief about whether I did this in the summer or the fall. I *entertain* the proposition that it was in the summer; I *consider* whether it is true; but, being utterly uncertain, I *suspend judgment* on it. I thus neither believe it nor *disbelieve* it, that is, believe it is false. My stance is one of *non-belief*. I need not try to force myself to resolve the question and judge the proposition either way. I might need to resolve it if something important turned on when I did the pruning; but here suspended judgment, with the resulting non-belief, is not uncomfortable.

As I think about cutting the vine, it occurs to me that in recalling that task, I am vividly imaging it. Here, I seem to be looking into my own consciousness, thus engaging in a kind of *introspection*. I can still see, in my mind's eye, the furry vine clinging to the tree, the ax, the sappy wound along the trunk where the vine was severed from it. I have turned my attention inward to my own imagery. The object of my attention, my own imaging of the scene, seems internal and is present to my consciousness, though *its* object is external and long gone by. But clearly, I believe that I am imaging the vine; and there is no apparent reason to doubt that I justifiedly believe this and know that it is so. This is a simple case of *self-knowledge*.

Reason and rational reflection

I now look back at the field and am struck by how perfectly rectangular it looks. If it is perfectly rectangular, then its corners are right angles. Here I believe something different in kind from the things cited so far: that if the field is rectangular, then its corners are right angles. This is a geometrical belief. I do not hold it on the same sort of basis I have for the other things I have mentioned believing. My conception of geometry as applied to ideal figures seems to be my basis. On that basis, my belief seems to be firmly justified and to constitute knowledge.

I can see that the spruce is taller than the maple, and that the maple is taller than the crab apple tree on the lawn closer by. I now realize that the spruce is taller than the crab apple. My underlying belief here is that if one thing is taller than a second and the second taller than a third, then the first is taller than the third. And, perhaps even more than the geometrical belief,

this abstract belief seems to arise simply from my grasp of the concepts in question, above all the concept of one thing's being taller than another.

Testimony

The season has been dry, and it now occurs to me that the roses will not flourish without a good deal of water. But this I do not believe simply on the basis of perception. One source from which I learned it is repeated *observation*. But there is another possible source: although much knowledge comes directly from our own experience, much also originates with *testimony* from others. I have received testimony as to where on the stem to trim off dead roses. If I did not learn about watering roses from my own experience, I could have learned the same things from testimony, just as I learned from a friend how far back to clip off dead roses.

To be sure, I need perception, such as hearing what I am told, to acquire knowledge on the basis of testimony, just as I needed perception to *learn* these things about roses on my own; and I need memory to *retain* them whatever their source. They are, however, generalizations and hence do not arise from perception in the direct and apparently simple way my visual beliefs do, or emerge from memory in the way my beliefs about past events I witnessed do. But do I not still justifiedly believe that the roses will not flourish without a lot of water? The commonsense view is that I both justifiedly believe and know this about roses, and that I can know it either through generalizing—a kind of reasoning—from my own observations, or from testimony, or from both.

Basic sources of belief, justification, and knowledge

The examples just given represent what philosophers have called perceptual, memorial, introspective, a priori, inductive, and testimony-based beliefs. The first four kinds are basic in epistemology. My belief that the glass is cold to the touch is *perceptual*, being based as it is on tactual perception. My belief that I cut the poison ivy vine from the spruce is *memorial*, since it is stored in my memory and held because of that fact. My belief that I am imagining a green field is called *introspective* because it is conceived as based on "looking within" (the etymological meaning of 'introspection'); but it could also be called simply *self-directed*: no "peering" within or special concentration is required. My belief that if the spruce is taller than the maple and the maple is taller than the crab apple then the spruce is taller than the crab apple is called *a priori* (meaning, roughly, based on what is "prior" to observational experience) because it apparently arises not from experience of how things actually behave, but simply in an *intuitive* way. It arises from a rational grasp of the key concepts one needs in order to have the belief, such as the concept of one thing's being taller than another.

By contrast, my belief that the roses will not grow well without abundant

water does not arise directly from one of the four basic sources just mentioned: perception, memory, introspection, and a priori intuition (reason, in one sense of the term). It is called *inductive* because it is formed (and held) on the basis of a generalization from something more basic, in this case what I learned from perceptual experiences with roses. Those experiences, apparently through my beliefs recording them, "lead into" the generalization about roses, to follow the etymological meaning of 'induction'. For instance, I remember numerous cases in which roses have faded when dry, and I eventually concluded that they need abundant water.

Each of the four basic kinds of belief I have described—perceptual, memorial, introspective, and a priori—is grounded in the source from which it arises. The nature of this grounding is explored in detail in Part One, which concerns perception, memory, consciousness, and reason. These sources are commonly taken to provide raw materials for inductive generalizations, as where observations and memories about roses yield a basis for generalizing about their needs.

Any of the beliefs we considered could instead have been grounded in testimony (the topic of Chapter 7), had I formed the beliefs on the basis of being given the same information by someone I trust. That person, however, would presumably have acquired it through one of these *other* sources (or ultimately through someone's having done so), and this makes testimony a different kind of source. This is why testimony is not a basic source of knowledge. It is still, however, incalculably important for human knowledge and unlimitedly broad. It can, for instance, justify a much wider range of propositions than perception can. We can credibly tell others virtually anything we know.

Three kinds of grounds of belief

Our examples illustrate not only grounding of beliefs *in* a source, such as perception or introspection, but also *how* they are grounded in these sources. There are at least three important kinds of grounding of beliefs—ways they are grounded. These are causal, justificational, and epistemic grounding. All three are important for many major epistemological questions.

Consider my belief that there is a green field before me. It is *causally grounded* in my experience of seeing the field because that experience produces or underlies the belief. It is *justificationally grounded* in that experience because the experience, or at least some element in the experience, justifies my belief. And it is *epistemically grounded* in the experience because in virtue of that experience my belief constitutes knowledge that there is a green field before me ('epistemic' comes from the Greek *episteme* meaning, roughly, 'knowledge'). These three kinds of grounding very often coincide (though Chapter 11 will describe important cases in which knowledge and justification do not). I will thus often speak simply of a belief as *grounded* in a source, such as visual experience, when what grounds the belief does so in all three ways.

Causal, justificational, and epistemic grounding each go with a very common kind of question about belief. Let me illustrate.

Causal grounding goes with 'Why do you believe that?' An answer to this, asked about my belief that there is a green field before me, would be that I see it. This is the normal kind of reply; but as far as mere causal production of beliefs goes, the answer could be brain manipulation or mere hypnotic suggestion. If, however, mere brain manipulation or mere hypnotic suggestion produces a belief, then the causal ground of the belief would not justify it. If, under hypnosis, I am told that someone dislikes me and as a result I believe this, the belief is not thereby justified.

Justificational grounding goes with such questions as 'What is your justification for believing that?' or 'What justifies you in thinking that?' or 'Why should I accept that?' ('Why do you believe that?' *can* be asked with this same justification-seeking force.) Again, I might answer that I see it. I might, however, have a justification (the situational kind) that, unlike seeing the truth in question, is not a *cause* of my believing it.

The justification I cite could also be the testimony of a credible good friend. It could be this even when, by a short circuit, brain manipulation does the causal work of *producing* my belief and leaves the testimony like a board that slides just beneath a roof beam but bears none of its weight. This shows that an element that provides situational justification for a belief may play no role in producing or supporting the belief, even if this element, like the auxiliary unstressed board, stands ready to play a supporting role if the belief is put under pressure by a challenge.

Epistemic grounding goes with 'How do you know that?' Once again, saying that I see it will commonly answer this. Here, however, it may be that a correct answer *must* cite something that is *also* a causal ground for the belief (a matter discussed in Chapter 10). Certainly a justificational ground need not be a ground of knowledge. One can justifiedly believe a proposition without knowing it.

Clearly, the same sorts of points can be made for the other five cases I have described: memorial beliefs are grounded in memory, self-directed ("introspective") beliefs in consciousness, inductively based beliefs in further, premise-beliefs that rest on experience, a priori beliefs in reason, and testimony-based beliefs in testimony.

Fallibility and skepticism

Even well-grounded beliefs can be mistaken. We can be deceived by our senses. We are fallible in perceptual matters, as in our memories, in our reasoning, and in other respects. One might now wonder, as skeptics do, whether we *know* even that it is improbable that our senses are now deceiving us. One might also wonder whether, when we take ourselves to see green grass, we are even *justified* in our belief that no such mistake has occurred.

Suppose that I am in an unfamiliar park. I might not know or even

justifiedly believe that artificial grass has not replaced the natural grass I take to be before me. (I may have heard of such substitutions and may have no good reason to believe this has not happened, though I do not consider the matter.) Am I justified in believing that there is green grass before me?

Suppose that I am not justified in believing there is green grass before me. If not, how can I be justified in believing what appear to be far less obvious truths, such as that my home is secure against the elements, my car safe to drive, and my food free of poison? And how can I know the many things I need to know in life, such as that my family and friends are trustworthy, that I can control my behavior and thus partly determine my future, and that the world we live in at least approximates the structured reality portrayed by common sense and science?

These are difficult and important questions. They indicate how insecure and disordered human life would be if we could not suppose that we possess justified beliefs and knowledge. We stake our lives every day on what we take ourselves to know. It would be unsettling to revise this stance and retreat to the view that at best we have justification to believe. But if we had to give up even this moderate view and to conclude, say, that what we believe is not even justified, we would face a crisis. Much later, in discussing skepticism, I will explore such questions at some length. Until then I will assume the commonsense view that beliefs with a basis like that of my belief that there is a green field before me are not only justified but also constitute knowledge.

Once we proceed on this commonsense assumption, it is easy to see that there are many different kinds of circumstances in which beliefs arise in such a way that they are apparently both justified and constitute knowledge. In considering this variety of circumstances yielding justification and knowledge, we can explore how beliefs are related to perception, memory, consciousness, reason, and testimony (the topics of Chapters 1–7).

Overview

There is a great deal more to be said about each of these sources of belief, justification, and knowledge and about how they ground what they do ground. The first seven chapters explore, and in some cases compare, the basic sources of belief, justification, and knowledge.

In the light of what those chapters show, we can discuss the development and structure of knowledge and justification (the task of Part Two). Much of what we believe does not come directly from perception, memory, introspection, or reflection of the kind appropriate to knowledge of such truths as those of elementary mathematics or those turning on our grasp of simple relations, for instance the proposition that if the spruce is taller than the maple, then the maple is shorter than the spruce, which we know by virtue of understanding the relations expressed by ‘taller’ and ‘shorter’. We must explore how inference and other developmental processes expand our body of knowledge and justified beliefs (this is the task of Chapter 8). Moreover,

once we think of a person as having the resulting complex body of knowledge and justified belief, we encounter the questions of what structure that large and intricate body has, and of how its structure is related to the amount and kind of knowledge and justification it contains. As we shall see in Part Two, these structural questions take us into an area where epistemology and the philosophy of mind often overlap.

On the basis of what Part One shows about sources of knowledge and justification and what Part Two shows about their development and structure, we can fruitfully proceed to consider more explicitly what knowledge and justification are and what kinds of things can be known (the task of Part Three). It is true that if we had no sense at all of what they are, we could not find the kinds of examples of them needed to explore their sources and their development and structure. If we do not have before us a wide range of examples of justification and knowledge, we lack the data appropriate to seeking a philosophically illuminating analysis of them. It is in the light of the examples and conclusions of Parts One and Two that Chapters 10 and 11 clarify the concept of knowledge, and, to a lesser extent, that of justification, in some detail.

With a conception of knowledge laid out, it is possible to explore the apparent extent of knowledge and justification in three major territories—the scientific, the ethical, and the religious. In exploring these domains, Chapter 12 applies some of the epistemological results of the earlier chapters. These chapters continue to assume the commonsense view that we have a great deal of knowledge and justification. If, however, skepticism is justified, then the commonsense assessment that the first twelve chapters make regarding the extent of knowledge and justification must be revised. Whether skepticism is justified is the focus of Chapters 13 and 14.

Along the way in all fourteen chapters, there is much to be learned about concepts that are important both in and outside epistemology, especially those of belief, causation, certainty, coherence, explanation, fallibility, illusion, inference, intellectual virtue, introspection, intuition, meaning, memory, rationality, reasoning, relativity, reliability, truth, and understanding. There are also numerous epistemological positions to be considered, sometimes in connection with historically influential philosophers. But the main focus will be on the major concepts and problems in the field, not on any particular philosopher or text. This may well be the best way to facilitate studying philosophers and epistemological texts; it will certainly simplify an already complex task.

Knowledge and justification are not only interesting in their own right as central epistemological topics; they also represent positive values in the life of every reasonable person. For all of us, there is much we want to know. We also care whether we are justified in what we believe—and whether others are justified in what they tell us. The study of epistemology can help in making this quest, even if it often does so indirectly. It can certainly help us assess how well we have done in the quest when we review our results.

Well-developed concepts of knowledge and justification can serve as ideals in human life. Positively, we can try to achieve knowledge and justification in relation to subjects that concern us. Negatively, we can refrain from forming beliefs where we think we lack justification, and we can avoid claiming knowledge where we think we can at best hypothesize. If we learn enough about knowledge and justification conceived philosophically, we can better search for them in matters that concern us and can better avoid the dangerous pitfalls that come from confusing mere impressions with justification or mere opinion with knowledge. This is not to say that epistemological knowledge can be guaranteed to yield new everyday knowledge. But the more we know about the constitution of knowledge and justification, the better we can build them through our own inquiries, and the less easily we will fall into the pervasive temptation to take an imitation to be the real thing.

Part One

Sources of justification, knowledge, and truth

1 Perception

Sensing, believing, and knowing

- **The elements and basic kinds of perception**
  Perceptual belief
  Perception, conception, and belief
  Propositional and objectual perception

- **Seeing and believing**
  Perceptually embedded beliefs
  Perception as a source of potential beliefs
  The perceptual hierarchy
  Simple, objectual, and propositional perception
  The informational character of perception

- **Perceptual justification and perceptual knowledge**
  Seeing and seeing as
  Perceptual content
  Seeing as and perceptual grounds of justification
  Seeing as a ground of perceptual knowledge

1 Perception

Sensing, believing, and knowing

As I look at the green field before me, I might believe not only that there is a green field there but also that I *see* one. And I do see one. I visually perceive it. Both beliefs, the belief that there is a green field there, and the self-referential belief that *I* see one, are grounded, causally, justificationally, and epistemically, in my perceptual experience. They are produced by that experience, justified by it, and constitute knowledge in virtue of it.

The same sort of thing holds for the other senses. Consider touch. I not only believe, through touch (as well as sight), that there is a glass here, I also feel its cold surface. Both beliefs—that there is a glass here and that it is cold—are grounded in my tactual experience. I could believe these things on the basis of someone's testimony. My beliefs would then have a quite different status. For instance, my belief that there is a glass here would not be a *perceptual belief*, but only a belief *about a perceptible*, that is, a perceivable object, the kind of thing that can be seen, touched, heard, smelled, or tasted. Through testimony we have beliefs about perceptibles we have never seen or experienced in any way.

My concern is not with the hodgepodge of beliefs that are simply about perceptibles, but with perception and perceptual beliefs. Perceptual beliefs are not simply beliefs about perceptibles; they are *beliefs grounded in perception*. We classify beliefs as perceptual by the nature of their roots, not by the color of their foliage; by their grounds, not their type of content. Those roots may be visual, auditory, and so forth for each perceptual mode. But vision and visual beliefs are an excellent basis for discussing perception, and I will concentrate on them and mention the other senses only when it adds clarity.

Perception is a source of knowledge and justification mainly by virtue of yielding beliefs that *constitute* knowledge or are justified. But we cannot hope to understand perceptual knowledge and justification simply by exploring those beliefs. We must also understand what perception is and *how* it yields beliefs. We can then begin to understand how it yields knowledge and justification or—sometimes—fails to yield them.

The elements and basic kinds of perception

There are apparently at least four elements in perception: (1) the perceiver, me; (2) the object, the field I see; (3) the sensory experience, say my visual experience of colors and shapes; and (4) the relation between the object and the subject, commonly taken to be a causal relation by which the object produces the sensory experience in the perceiver. To see the field is apparently to have a certain sensory experience as a result of the impact of the field on our vision.

Some accounts of perception add to the four items on this list; others subtract from it. To understand perception we must consider both kinds of account and how these elements are to be conceived in relation to one another. But first, it is essential to explore examples of perception.

There are several quite different ways to speak of perception. Each corresponds to a different way of perceptually responding to experience. We often speak simply of what people perceive, for instance see. We also speak of what they perceive the object to be, and we commonly talk of facts they know through perception, such as that the grass is long. Visual perception most readily illustrates this, so let us start there.

I see, hence perceive, the green field. Second, speaking in a less familiar way, I see it *to be* rectangular. Thus, I might say that I know it looks irregular from the nearby hill, but from the air you can see it to be perfectly rectangular. Third, I see *that it is* rectangular. Perception is common to all three cases. Seeing, which is a paradigm perception, is central in each.

The first case is one of *simple perception*, perception taken by itself (here, visual perception). I see the field, and this experience is the visual parallel of hearing a bird (an auditory experience), touching a glass (a tactual experience), smelling roses (an olfactory experience), and tasting mint (a gustatory experience). If the first case is simply *perceiving of* some object, the second is a case of *perceiving to be*, as it is seeing something to be so: I do not just see the field, as when I drive by at high speed and do not even realize what is in my peripheral vision; rather, I see the field to be rectangular. The third case is one of *perceiving that*; it is seeing that a particular thing is so, namely that the field is rectangular.

These cases represent three kinds, or *modes*, of perception. Perception of the simplest kind (or in the simplest mode), such as seeing, occurs in all three; but, especially because of their relation to knowledge and justified belief, they are significantly different. We can best understand these three kinds (or modes) of perception if we first focus on their relation to belief.

Perceptual belief

The last two cases—perceiving *that*, and perceiving *to be*—are different from the first—perceiving *of*—in implying corresponding kinds of beliefs: seeing that the field is rectangular implies believing that it is, and seeing it

to be green implies believing it to be green. If we consider how both kinds of beliefs—beliefs *that* something is so and beliefs *of* (hence *about*) something—are related to perception, we can begin to understand how perception occurs in all three cases, the simple and the more complex. In my second and third examples of perception, visual perception (seeing) issues in beliefs that are grounded in seeing and can thereby constitute visual knowledge, such as knowing that the field is green.[1]

In our example of simple perception, my just seeing the field provides a basis for both kinds of beliefs. It does this even if, because my mind is entirely occupied with what I am hearing on the radio as I glance over the field, no belief about the field actually arises in me. The visual experience is, in this instance, like a foundation that has nothing built on it but is ready to support a structure. If, for example, someone were to ask if the field has shrubbery, then given the lilacs prominent in one place, I might immediately form the belief that it does and assent. This belief is visually grounded; it comes *from* my seeing the field though it did not initially come *with* it. When visual experiences do produce beliefs, as they usually do, what kinds of beliefs are these, and how are they specifically perceptual?

Many of my beliefs arising through perception correspond to perception *that*, say to seeing that the lilacs are blooming. I believe that the field is lighter green toward its borders, that it is rectangular in shape, and that it has many ruts. But I may also have various beliefs about it that are of the second kind: they correspond to perception *to be*, for instance to seeing something to be a certain color. Thus, I believe the field to be green, to be rectangular, and so on. The difference between these two kinds of belief is significant. As we shall shortly see, it corresponds first of all to two distinct ways in which we are related to the objects we perceive and, second, to two different ways of assessing the truth of what, on the basis of our perceptions, we believe.

The first kind of belief just described is the kind people usually think of when they consider beliefs: it is called *propositional*, as it is generally considered a case of believing a proposition—say, *that* the field is rectangular. The belief is thus true or false depending on whether the proposition in question—here *that the field is rectangular*—is true or false. In holding the belief, moreover, in some way I think of what I see as a field which is rectangular: in believing that the field is rectangular, I *conceive* what I take to be rectangular *as* a field.

The second kind of belief might be called *objectual*: it is a belief regarding an object, say the field, with which the belief is actually connected. This is an object *of* (or about) which I believe something, say that it is rectangular. If I believe the field to be rectangular, there really is such an object, and I have a certain relation to it. A special feature of this relation is that there is no particular proposition I must believe about the field. To see that there is no particular proposition, notice that in holding this objectual belief I need not think of what I see *as* a field. I might mistakenly take it to be (for instance) a lawn or a grasslike artificial turf, yet still believe it to be rectangular. I might

think of it just in terms of what I believe it to be and not in terms of anything else.

Thus, although there is *some* property I must take the field to have—corresponding to what I believe it to *be*—there is no other particular way I must think of it. With objectual belief, then, there is no particular notion, no specific conceptual "handle," that must yield the subject of any proposition I believe about the object: I do not have to believe that the *field* is green, that the *grass* is green, or any such thing. Perception leaves us vast latitude as to what we learn from it. People differ greatly in the beliefs they form about the very same things they see.[2]

The concept of objectual perception, then, is very permissive about what one believes about the object perceived. This is one reason why it leaves so much space for imagination and learning—a space often filled by the formation of propositional beliefs, each capturing a different aspect of what is perceived, say that the field is richly green, that it is windblown, and that it ends at a treeline.

A different example may bring these points out further. After seeing a distant flare and coming to believe, of something blurry and far away, that it glowed, one might ask, 'What on Earth was it that glowed?' Before we can believe the proposition that a *flare* glowed, we may have to think about where we are, the movement and fading of the glow, and so forth. The objectual belief is a guide by which we may arrive at propositional beliefs and propositional knowledge.

Perception, conception, and belief

The same kind of example can be used to illustrate how belief depends on our conceptual resources in a way that perception does not. Suppose I had grown up in the desert and somehow failed to acquire the concept of a field. I could nonetheless see the green field, and from a purely visual point of view it might *look* the same to me as it does now. I could also believe, regarding the field I see—and perhaps conceive as sand artificially covered with something green—that it is rectangular. But I could not believe that the *field* is rectangular. This propositional belief as it were portrays what I see *as* a field in a way that requires my having a concept of one.

There is a connection here between thought and language (or at least conceptualization). If I believe (think) that the field is rectangular, or even simply have the thought that it is, I should be able to *say* that it is and to know what I am talking about. But if I had no concept of a field, then in saying this I would not know what I am talking about.[3] Similarly, a two year old, say, Susie, who has no notion of a tachistoscope, can, upon seeing one and hearing its fan, believe it to be making noise; but she cannot believe specifically that the tachistoscope is making noise. Her propositional belief, if any, would be, say, that the thing on the table is making noise. Since this is true, what she believes is true and she may know this truth, but she need not know much

about the object this truth concerns: in a way, she does not know what it is she has this true belief *about*.

The general lesson here is important. A basic mode of learning about objects is to find out truths about them in this elementary way: we get a handle on them through perceptually discriminating some of their properties; we form objectual (and other) beliefs about them from different perspectives; and (often) we finally reach an adequate concept of what they are. From the properties I believe the flare in the distance to have, I finally figure out that it is a flare that has them. This suggests that there is at least one respect in which our knowledge of (perceptible) *properties* is more basic than our knowledge of the substances that have them; but whether that is so is a question I cannot pursue here.

Unlike propositional beliefs, objectual beliefs have a significant degree of indefiniteness in virtue of which it can be misleading simply to call them true or false; they are accurate or inaccurate, depending on whether what one believes of the object (such as that it is rectangular) is or is not *true of* it. Recall Susie. If she attributes noise-making to the tachistoscope, she truly believes, *of* it, that it is making noise. She is, then, *right about it*. But this holds even if she has no specific concept of what it is that is making the noise. If we say unqualifiedly that her belief about it is true, we invite the question 'What belief?' and the expectation that the answer will specify a particular proposition, say that the tachistoscope is making noise. But it need not, and we might be unable to find any proposition that she does believe about it. She can be right about something without knowing or even having any conception of what kind of thing it is that she is right about.

Knowledge is often partial in this way. Still, once we get the kind of epistemic handle on something that objectual belief can provide, we can usually use that to learn more about it.[4] Suppose I see a dog's tail projecting from under a bed and do not recognize it as such. If I believe it to be a slender furry thing, I have a place to start in finding out what else it is. I will, moreover, be disposed to form such beliefs as that there is a slender furry thing there. I will also have justification for them. But I need not form them, particularly if my attention quickly turns elsewhere.

Propositional and objectual perception

Corresponding to the two kinds of beliefs I have described are two ways of talking about perception. I see *that* the field is rectangular. This is (visual) *propositional perception*: perceiving *that*. I also see it *to be* rectangular. This is (visual) *objectual perception*: perceiving *to be*. The same distinction apparently applies to hearing and touch. Perhaps, for example, I can hear that a piano is out of tune by hearing its sour notes, as opposed to hearing the tuner say it needs tuning. As for taste and smell, we speak as if they yielded only simple perception: we talk of smelling mint in the iced tea, but not of smelling that it is minty or smelling it to be minty. Such talk is, however,

intelligible on the model of seeing that something is so and seeing it to be so. We may thus take the distinction between perceiving *that* and perceiving *to be* to apply in principle to all the senses.

It is useful to think of perceptual beliefs as *embedded* in the corresponding propositional or objectual perception, roughly in the sense that they are integrally tied to perceiving of that kind and derive their character and perhaps their authority from their perceptual grounding. Take propositional belief first. My belief that the field is rectangular is embedded in my seeing that it is, and Susie's believing the tachistoscope to be making noise is embedded in her hearing it to be doing so. In each case, the belief is an element in perception of the corresponding kind. These kinds of perception might therefore be called *cognitive*, since belief is a cognitive attitude: roughly the kind having a proposition (something true or false) as its object.[5] The object of the belief that the field is rectangular is the specific proposition that the field is rectangular, which is true or false.

Now consider objectual perceptual beliefs. If believing the tachistoscope *to be* making noise has a propositional object, that object may be plausibly taken to be some proposition or other to the effect that it is making noise, which (though left unspecified by the ascription of the belief) is also true or false. But some objectual perceptions may also be plausibly conceived as simply attributions of a perceptible property to the thing perceived; here the embedded objectual belief is *true of* the object rather than simply true. A tiny, prelingual child might see the liquid offered to it to be milk yet not believe (or disbelieve) the proposition that it is milk. In this respect, belief is unlike attitudes of approval or admiration or indignation, which are evaluated not as true or false but rather as, say, appropriate or inappropriate.[6]

Both propositional and objectual beliefs are grounded in simple perception. If I do not see a thing at all, I do not see that it has any particular property and I do not see it to be anything. Depending on whether perceptual beliefs are propositional or objectual, they may differ in the kind of knowledge they give us. Propositional perception yields knowledge both of *what* it is that we perceive and of some *property* of it, for instance of the *field's* being *rectangular*. Objectual perception may, in special cases, give us knowledge only of a property of what we perceive, say of its being green, when we do not know what it is or have any belief as to what it is.

In objectual perception, we are, to be sure, in a good position to come to know *something* or other about the object, say that it is a green expanse. Objectual perception may thus give us information not only about objects of which we have a definite conception, such as home furnishings, but also about utterly unfamiliar objects of which we have at most a very general conception, say 'that noisy thing'. This is important. We could not learn as readily from perception if it gave us information only about objects we conceive in the specific ways in which we conceive most of the familiar things we see, hear, touch, taste, and smell.[7]

Seeing and believing

Both propositional and objectual perceptual beliefs are commonly grounded in perception in a way that apparently connects us with the outside world and assures their truth. For instance, my visual belief that the field is rectangular is so grounded in my seeing the field that I veridically (truly) see that it is rectangular; my tactually believing the glass to be cold is so grounded in my feeling it that I veridically feel it to be cold. Let us explore the relation between perception and belief.

Perceptually embedded beliefs

Must beliefs grounded in seeing be true? Admittedly, I might visually (or tactually) believe that something is rectangular under conditions poor for judging it. Compare viewing a straight stick half submerged in water (it will look bent). My visually grounded belief might then be mistaken. But such a mistaken belief is not *embedded* in propositional perception that the stick is bent—that proposition is false and hence is *not* something one sees is so (or to be so). The belief is merely *produced* by some element in the simple perception of the stick: I see the stick in the water, and the operation of reflected light causes me to have the illusion of a bent stick. I thus do not see that the stick is bent: my genuine perception is of it, but not of its curvature. Seeing that curvature or seeing that the stick is bent would entail that it *is* bent, which is false. If the stick is not bent, I cannot see that it is.

As this suggests, there is something special about both perceiving *that* and perceiving *to be*. They are *veridical experiences*, that is, they imply truth. Specifically, if I see that the field is rectangular, or even just see it to be rectangular, then it truly is rectangular. Thus, when I simply see the rectangularity of the field, if I acquire the corresponding embedded perceptual beliefs—if I believe that it is rectangular when I see that it is, or believe it to be rectangular when I see it to be—then I am correct in so believing.

Perceiving *that* and perceiving *to be*, then, imply (truly) believing something about the object perceived—and so are *factive*. Does simple perception, perception *of* something, which is required for either of these more complex kinds of perception, also imply true belief? Very commonly, simple perception does imply truly believing something about the object perceived. If I hear a car go by, I commonly believe a car is passing. But could I not hear it, but be so occupied with my reading that I form no belief about it? Let us explore this.

Perception as a source of potential beliefs

As is suggested by the case of perception overshadowed by preoccupation with reading, there is reason to doubt that simple perceiving *must* produce any belief at all. Moreover, it commonly does not produce beliefs even of

what *would* be readily believed if the question arose. Suppose I am looking appreciatively at a beautiful rug. Must I believe that it is not producing yellow smoke, plain though this fact is? I think not; there seems to be a natural economy of nature—perhaps explainable on an evolutionary basis—that prevents our minds from being cluttered with the innumerable beliefs we would have if we formed one for each fact we can see to be the case.

This line of thought may seem to fly in the face of the adage that seeing is believing. But properly understood, that may apply just to propositional or objectual seeing. In those cases, perception plainly does entail beliefs. Seeing that golf ball-size hail is falling *is* (in the sense that it entails) believing it.[8] This fact, however, is not only perceptible; it is striking.

In any event, could I see the field and believe nothing regarding it? Must I not see it to be something or other, say green? And if so, would I not believe, of it, *something* that is true of it, even if only that it is a green object some distance away? Consider a different example.

Imagine that we are talking excitedly and a bird flies quickly across my path. Could I see it, yet form no beliefs about it? There may be no clearly correct answer. For one thing, although there is much we *can* confidently say about seeing and believing, 'seeing' and 'believing' are, like most philosophically interesting terms, not precise. They have an element of vagueness. No standard dictionary definition or authoritative statement can be expected either to tell us precisely what they mean or, especially, to settle every question about when they do and do not apply.[9] Still, we should be wary of concluding that vagueness makes any significant philosophical question unanswerable. How, then, should we answer the question whether seeing entails believing?

A negative response might be supported as follows. Suppose I merely see the bird but pay no *attention* to it because I am utterly intent on our conversation. Why must I form any belief about the bird? Granted, if someone later asks if I saw a blue bird, I may assent, thereby indicating a belief that the bird *was* blue. But this belief is not perceptual: it is about a perceptible and indeed has visual content, but it is not grounded in seeing. Moreover, it may have been formed only when I recalled my visual experience of the bird. Recalling that experience in such a context may *produce* a belief *about* the thing I saw even if my original experience *of* the thing did not. For plainly a recollected sensory experience can produce beliefs about the object that caused it, especially when I have reason to gain information about that object. Perhaps one notices something in one's recollected image of the bird, an image merely recorded in the original experience, but one formed no belief about the bird. Granted, perception must produce a sensory experience, such as an image, and granted such an image—and even a recollection of it—is raw material for beliefs; it does not follow that perception must produce beliefs.

It might be objected that genuinely seeing an object must produce beliefs, even if we are not *conscious* of its doing so. How else can perception guide our behavior, as it does when, on seeing a log in our path, we step over it?

One answer is that not everything we see, including the bird that flies by as I concentrate on something else, demands or even evokes a cognitive response, particularly one entailing belief-formation. If I am cataloguing local birds, the situation is different. But when an unobtrusive object we see—as opposed to one blocking our path—has no particular relation to what we are doing, perhaps our visual impressions of it are simply a *basis* for forming beliefs about it should the situation call for it, and it need not produce any belief if our concerns and the direction of our attention give the object no significance.

Despite the complexity I am pointing to in the relation between seeing and believing, clearly we may hold what is epistemologically most important here. Suppose I can see a bird without believing anything about (or of) it. Still, when I do see one, I *can* see it to be something or other, and my perceptual circumstances are such that I might readily both come to believe something about it *and* see that to be true of it. Imagine that someone suddenly interrupts a conversation to say, 'Look at that bird!' If I see it, I am in a position to form some belief about it, if only that it is swift, though I need not actually form any belief about it, at least not one I am conscious of.

To see these points more concretely, imagine I am alone and see the bird in the distance for just a second, mistakenly taking it to be a speck of ash. If there is not too much color distortion, I may still both know and justifiedly believe it to be dark. Granted, I would misdescribe it, and I might falsely believe that it is a speck of ash. But I could still know something about it, and I might point the bird out under the misleading but true description, 'that dark thing'. The bird *is* the thing I point at; and I can see, know, and justifiedly believe that there is a dark thing there.

My perception of the bird, then, gives me a ready *basis* for some knowledge and justification, even if the perception occurs in a way that does not cause me to believe that there is, say, a *bird* before me and so does not give me *actual* knowledge of it. Seeing *is* virtual believing, or at least potential believing. A similar point holds for simple perception in the other senses, though some, such as smell, are in general less richly informative than sight.[10]

The perceptual hierarchy

Our discussion seems to show that simple perceiving need not produce belief, and objectual perceiving need not always yield propositional perceiving. Still, this third kind of perception is clearly not possible without the first and, I think, the second as well. I certainly cannot see *that* the bird is anything if I do not see it at all; and I must also see it in order to see it *to be* something, say a speck of blue. Thus, simple perceiving is fundamental: it is required for objectual and propositional perceiving, yet does not clearly entail either. If, for instance, you do not perceive in the simple mode, say see a blue speck, you do not perceive in the other two modes either, say see a speck to be blue or see that it is blue. And as objectual perceiving seems possible without

propositional perceiving, but not conversely, the former seems basic relative to the latter.

Simple, objectual, and propositional perception

We have, then, a perceptual hierarchy: propositional perceiving depends on objectual perceiving, which in turn depends on simple perceiving. Simple perceiving is basic, and it commonly yields, even if it need not always yield, objectual perceiving, which, in turn, commonly yields, even if it need not always yield, propositional perceiving. Simple perceiving, such as just seeing a green field, may apparently occur without either of the other two kinds, but seeing something *to be* anything at all, such as rectangular, requires seeing it; and seeing *that* it is something in particular, say green, requires both seeing it to be something and, of course, seeing it.

Thus, even if simple perception does not always produce at least one true belief, it characteristically does position us to form any number of true beliefs. It gives us cognitive *access* to perceptual information, perhaps even *records* that information in some sense, whether or not we register the information conceptually by forming perceptual beliefs of either kind.

The informational character of perception

As this suggests, perception by its very nature is *informational*; it might even be understood as equivalent to a kind—a sensory kind—of receipt of information about the object perceived.[11] The point here is that not all perceptually given information is *propositional* or even conceptualized. This is why we do not receive or store all of it in the contents of our beliefs. Perceptual content—conceived as the content of a simple perception—is at least in part determined by the properties we are sensorily conscious of in having that experience; it is not equivalent to the content of the perceptual belief(s) that experience may produce.

Some of the information perception yields is imagistic. Indeed, we may think of all the senses as capable of yielding images or, for the non-visual senses, at least of yielding the non-visual counterparts of images—*percepts*, to use a technical term for such elements in perceptual experience occurring in any sensory mode, whether visual or auditory or of some other kind. It is in these sensory impressions that the bulk of perceptual information apparently resides. This point explains the plausibility of the idea that a picture is worth a thousand words—which is not to deny that, for some purposes, some words are worth a thousand pictures. A single report of smoke may avert a catastrophe; a single promise may alter a million lives.

It is in part because perception is so richly informative that it normally gives us not only imagistic information but also situational justification. Even if I could be so lost in conversation that I form no belief about the passing bird, I am, as I see it pass, normally justified in believing something about it,

concerning its perceptible properties, for instance that it glides.[12] There may perhaps be nothing highly specific that I am justified in believing about it, say that it is a cardinal or that its wingspan is ten inches, but if I really see it, as opposed to its merely causing in me a visual impression too indistinct to qualify me as seeing it, then there is something or other that I may justifiably believe about it.

When we have a clear perception of something, it is even easier to have perceptual justification for believing a proposition about it without actually believing it. Just by taking stock of the size of the field in clear view before me, I am justified in believing that it has more than 289 blades of grass; but I do not ordinarily believe—or disbelieve—any such thing about grassy fields I see. It was only when I sought a philosophical example about perception and belief, and then arbitrarily chose the proposition that the field has more than 289 blades of grass, that I came to believe this proposition. Again, I was justified in believing the proposition before I actually did believe it.

Perceptual justification and perceptual knowledge

What is it that explains why seeing the bird or the field justifies us in believing something about what we see, that is, gives us situational justification for such a belief? And does the same thing explain why seeing something enables us to know various facts about it?

Seeing and seeing as

One possible answer is that if we see something at all, say a bird, we see it *as* something, for instance black or large or swift, and we are justified in believing it to be what we see it as being. The idea is that all seeing and perhaps all perceiving is *aspectual perception* of a kind that confers justification. We see things by seeing their properties or aspects, for instance their colors or their front sides, and we are justified in taking them to have the properties or aspects we see them as having.

Let us not go too fast. Consider two points, one concerning the nature of seeing *as*, the other its relation to justification.

First, might not the sort of distinction we have observed between situational and belief justification apply to seeing itself? Specifically, might not my seeing the bird imply that I am only in a *position* to see it *as* something, and not that I *do* see it as something? It is true that when we see something, we see it *by* seeing some property or aspect of it; but it does not follow that we see it as *having* this property or aspect. I might see a van Gogh painting by its colors, shapes, and distinctive brush strokes, but not see it as having them because my visual experience is dominated by the painting as a whole. Someone might reply that if I see it by those properties, I am disposed to believe it has them and so must see it as having them; but this disposition

implies at most a readiness to see it as having them. There may, to be sure, be a sense in which if we see something aright, for example see a van Gogh with recognition of it as his, then we must see it as what we recognize it to be.

Seeing *as* can also be a matter of conceptualization—roughly, *conceiving as*. But this is different from perceptual seeing as. The distinction between perceptual seeing *as* and perceptual seeing *by* remains. Seeing *by* is causal and discriminative but not necessarily ascriptive or, especially, conceptual. Seeing *as,* though also causal, is often ascriptive and commonly conceptual. We see faces by seeing (for example) the distinctive shape of the eyes and mouth, but need not ascribe those to those we see or conceptualize these properties. But if we see a painting as blurry, we commonly ascribe that property to it and may conceptualize the painting as blurry.

Second, suppose that seeing the bird did imply (visually) seeing it *as* something. Clearly, this need not be something one is justified in believing it to be (and perhaps it need not be something one *does* believe it to be). Charles, our biased birdwatcher, might erroneously see a plainly black bird as blue, simply because he so loves birds of blue color and so dislikes black birds that (as he himself knows) his vision plays tricks on him when he is bird-watching. He might then not be justified in believing that the bird is blue.

Assume for the sake of argument that seeing implies seeing *as* and that typically, seeing as implies at least objectually believing something or other about the thing seen. Still, seeing an object as having a certain property—say, a stick in the water as bent—does not entail that it has the property. Nor does it always give one (overall) situational justification for believing it to have that property.

Perceptual content

It is natural to think of perception as in some way *representational*. If we see things by seeing their properties, for instance, then our perceptual experience in some way represents the object as having them. If perceiving entailed believing, we could perhaps take it to have the same content of the entailed belief(s). But (simple) perception apparently does not entail believing, so this conception of its content is mistaken. For propositional and objectual perception, however, we might plausibly say something like this: the content of my perception that *p* includes both the proposition that *p* (hence also the content of that proposition) and also the content of my objectual perception of the thing in question; that content includes the properties I perceive the thing to have.

If we seek a broad notion of perceptual content for simple perception, we might say that all the properties represented in a perceptual experience constitute its content. Then, for greater specificity, we might call the totality of perceptually represented properties the *property content*. These include properties an object is seen *as* having.[13] They apparently also determine "what it

is like" to perceive the object, say a squirrel in a tree. In seeing it, one's visual field is determined mainly by the grey, the distinctive furry shape, and the arboreal background.

For propositional and objectual perception, we might call the property-ascriptive propositions that the perceiver perceptually believes on the basis of the perceptual experience their *doxastic propositional content.* If we want to capture all the propositions that one might justifiedly believe (and know) on the basis of the perception, we might speak of its *total propositional content*. This would include such propositions as that the squirrel is crouching, has a nut in its mouth, is in sunlight, and many more that need not be believed as a result of simply seeing the animal.[14]

Seeing as and perceptual grounds of justification

Whether or not seeing always implies seeing *as*, it does have property content and normally puts one in a position to form at least one justified belief about the object seen. Suppose I see the bird so briefly and distractedly that I do not see it as anything in particular; still, my visual impression of it has some feature or other by which I am justified in believing something of the bird, if only that it is a moving thing. Even Charles would be justified in believing something like this. His tendency to see black birds as blue is irrelevant to his perception of movement and does not affect his justification for believing such moving objects to be in motion.

Suppose, however, that for hours Charles had been hallucinating all manner of unreal things, and he knows this. Then he might not be justified in taking the bird he sees to be *anything* real, even though it is real. For as a rational person in this position he should see that if his belief is true, it may well be true only in the way a lucky guess is. Thus, the best conclusion here—and I suggest that this is an important justification principle concerning perception—is that normally, seeing an object gives one situational justification for believing something or other about it.

More broadly, it is very plausible to hold that *the evidence of the senses*—including above all the sensory experiences characteristic of perception—normally provides justification for beliefs with content appropriate to that evidence. If your experience is of a green expanse, you are justified in believing there is something green before you; if it is of something cool in your hand, you are justified in believing there is something cool in your hand; and so on.

One might also say something slightly different, in a terminology that is from some points of view preferable: seeing an object (always) gives one *prima facie justification* for believing something or other about it. Prima facie justification is roughly justification that prevails unless *defeated.* The two main kinds of defeater are such overriding factors as a strong justification for believing something to the contrary and such undermining (or undercutting) factors as my knowledge that I have been hallucinating and at present cannot

trust my senses. Overriders defeat prima facie justification by justifying an incompatible proposition instead; undermining defeaters simply prevent the would-be justification from succeeding. If I see a green field, I have a justification for believing it to be green; but I may not be justified, *overall*, in believing this if credible friends give me compelling reason to believe that despite appearances the field is entirely covered by blue grass, or that I am not seeing a field at all but hallucinating one.[15] In the former case, my justification is defeated by my acquiring better justification for a contrary proposition; in the latter, my visual justification is reduced below the threshold of success. If it is not eliminated, it is too weak to license saying I am justified in believing the proposition.

If seeing is typical of perception in (normally) putting us in a position to form at least one justified belief about the object seen, then perception in general normally gives us at least situational justification. This is roughly justification *for holding* a belief of the proposition for which we have the justification. As our examples show, however, it does not follow that every perceptual belief is justified. Far from it. Some perceptual beliefs, such as perceptual beliefs that are evidentially undermined by one's having formed similar beliefs based on hallucinations, are not. As with the biased bird-watcher, belief can be grounded in perception under conditions that prevent its being justified by that grounding.

Nevertheless, there is a simple principle of justification we can see to be plausible despite all these complexities: normally, a visual belief that is embedded in seeing that something is so or in seeing it *to be* so is justified (and it is always prima facie justified). If we see that an object has a property (say, that a field is rectangular) and, in virtue of seeing that it has that property (say, is rectangular), believe that it does, then (normally) we justifiedly believe that it does. Call this *the visual justification principle*, since it applies to cases of belief based on seeing that what is believed is true (or seeing it to be true).

I say *normally* (and that the justification is prima facie) because even here one's justification can be defeated. Thus, Charles might see that a bird is blue and believe on this basis that it is, yet realize that all morning he has been seeing black birds as dark blue and thus mistaking the black ones for the blue ones. Until he verifies his first impression, then, he does not justifiedly believe that the bird is blue, even though it in fact is. (We could say that he has *some* justification for believing this, yet better justification for not believing it; but to simplify matters I am ignoring degrees of justification.) He does indeed see a bird and may justifiedly believe that, but his belief that the bird is blue is not justified.

Suppose, on the other hand, that Charles has no idea that he has been hallucinating. Then, even when he does hallucinate a blue bird, he may be justified in believing that there is a blue bird before him. This suggests a related principle of justification, one that applies to visual experience whether it is a case of seeing or merely of visual hallucination: *When, on the basis of an apparently normal visual experience (such as the sort we have in seeing a bird*

nearby), one believes something of the kind the experience seems to show (for instance that the bird is blue), normally this belief is justified. Call this the *visual experience principle*, since it applies to cases in which one has a belief based on visual experience even if not an experience of actually seeing (the veridical kind). The visual principle takes us from seeing (vision) to justification; the visual experience principle takes us from visual experience—conceived as apparent seeing—to justification. The latter is wider: it indicates that visual experience can justify a huge range of beliefs, not just a belief to the effect that an object in fact has a property one sees it to have.

Similar principles can be formulated for all of the other senses, though the formulations will not be as natural. If, for example, you hear a note to be flat and on that basis believe that it is flat, normally your belief is justified. It is grounded in a veridical perception in which you have discriminated the flatness you believe the note has. And suppose, by contrast, that in what clearly seem to be everyday circumstances you have an utterly normal-seeming auditory hallucination of a flat note. If that experience makes it seem clear that you are hearing a flat note, then if you believe on the basis of the experience that this is a flat note, normally your belief would be justified. You have no reason to suspect hallucination, and the justification of your belief that the note is flat piggybacks, as it were, on the principle that normally applies to veridical beliefs.[16]

Seeing as a ground of perceptual knowledge

Some of what holds for the justification of perceptual beliefs also applies to perceptual knowledge. Seeing the green field, for instance, normally yields knowledge about the field as well as justified belief about it. This suggests another visual principle, a visual knowledge principle. It might be called an *epistemic principle*, since it states a condition for the visual generation of *knowledge*: *At least normally, if we see that a thing (such as a field) has a property (say is rectangular), we (visually) know that it has it*. A parallel principle holds for objectual seeing: At least normally, if I see something to have a property (say to be rectangular), I know it to have the property.

There are, however, special circumstances that explain why these epistemic principles may have to be restricted to "normal" cases. It may be possible to see that something is so, believe on that basis that it is, and yet not know that it is. Charles's case seems to show this. For if, in the kind of circumstances he is in, he often takes a black bird to be blue, then even if he sees that a certain blue bird is blue and, on that basis, believes it is blue, he apparently does not know that it is.[17] He might as well have been wrong, one wants to say; he is just lucky that this time his belief is true and he was not hallucinating. As he has no reason to think he has been hallucinating, and does not realize he has been, one cannot fault him for holding the belief that the bird is blue or regard the belief as inappropriate to his situation. Still, knowledge apparently needs better grounding than is provided by his blameless good fortune. This

kind of case has led some philosophers to maintain that when we know that something is so, our being right is not *accidental*.

There is an important difference here between knowledge and justification. Take knowledge first. If Charles is making errors like this, then even if he has no idea that he is and no reason to suspect he is, he does not know that the bird he believes to be blue is blue. But even if he has no idea that he is making errors, or any reason to suspect he is, he *may* still justifiedly believe that the bird is blue. The main difference between knowledge and justification here may be this: he can have a true belief that does not constitute knowledge because there is something wrong for which he is in no way criticizable (his errors might arise from a handicap which he has no reason to suspect, such as sudden color blindness); but he cannot have a true yet unjustified belief without being in some way criticizable. The standards for knowledge, one might say, permit fewer unsuspected weaknesses in discriminating the truth than those for justification, if the standards for knowledge permit any at all.

This difference between knowledge and justification must be reflected in the kinds of principles that indicate how justification, as opposed to knowledge, is generated. Justification principles need not imply that the relevant basis of a belief's justification assures its truth; but since a false belief cannot constitute knowledge, epistemic principles (knowledge principles) cannot capture elements that generate knowledge unless they rule out factors that might produce a false belief. A ground of knowledge must, in *some* way, suffice for the truth of the proposition known; a ground of justification must, in some way, *count toward* the truth of the proposition one is justified in believing, but need not rule out its falsehood.

On the basis of what we see, hear, feel, smell, and taste, we have a great many beliefs, propositional and objectual. There is apparently no good reason to doubt that these perceptual beliefs are commonly justified or that, quite often, they are true and constitute knowledge. But to see that perception is a basis of justification and knowledge is to go only part way toward understanding what perception, justification, and knowledge are. Here the main question is what constitutes perception, philosophically speaking. Until we have a good understanding of what it is, we cannot see in detail how perception grounds belief, justification, and knowledge. These problems cannot be fully resolved in this book, but we can achieve partial resolutions. I want to discuss (further) what perception is first and, later, to illustrate in new ways how it grounds what it does. The next chapter, then—also concentrating on vision—will start by considering some of the major theories of the nature of perception.

Notes

1 Perceiving *of*, perceiving *to be*, and perceiving *that* may also be called perception of, perception to be, and perception that, respectively; but the

second expression is not common, and in that case at least, the *-ing* form usually better expresses what is intended.

2 A related way to see the difference between objectual and propositional beliefs is this. If I believe something to have a property, say a British Airways plane to be a Boeing 777, then this same belief can be ascribed to me using any correct description of that plane, say, as the most traveled plane in the British Airways fleet: to say I believe BA's most traveled plane to be a 777 is to ascribe the same belief to me. This holds even if I do not believe it meets that description—and it can hold even when I cannot understand the description, as a child who believes a tachistoscope to be making noise cannot understand 'tachistoscope'. By contrast, if I have a propositional belief, say that the United Airlines plane on the runway is the most traveled in its fleet, this ascription cannot be truly made using just any correct description of that plane, say the plane on which a baby was delivered on Christmas Day, 2001. I may have no inkling of that fact—or may mistakenly think it holds for a BA plane. A rough way to put part of the point here is to say that propositional beliefs about things are about them *under a description or name*, and objectual beliefs about things are not (even if the believer could describe them in terms of a property they are believed to have, such as being noisy). It is in part because we need not conceptualize things—as by thinking of them under a description—in order to have objectual beliefs about them that those beliefs are apparently more basic than propositional ones.

3 In terminology common in epistemology, objectual belief is *de re*—of the thing—whereas propositional belief is *de dicto*—of the proposition—and I am similarly distinguishing between objectual and propositional perception. The objectual cases, unlike the propositional ones, require no particular concept of the thing perceived. To be sure, those who do have the concept of a field and know that I believe it to be rectangular may say, 'He believes the field is rectangular', meaning that I believe it *to be* rectangular. English idiom is often permissive in this way, and in everyday life nothing need turn on the difference. Moreover, some philosophers have held that a thing, such as a field, can be a constituent in a proposition—in which case it might be considered a kind of content of a belief of that proposition—and this might provide a basis for saying that the two belief ascriptions may be properly interchangeable. I am ignoring that controversial and uncommon conception of a proposition. For detailed discussion of the extent to which perception is conceptual and of how it yields perceptual beliefs, see Michael Pendelbury, 'Sensibility and Understanding in Perceptual Judgments', *South African Journal of Philosophy* 18, 4 (1999), 356–69.

4 It may be best to leave open here that Susie could, at least for a moment, believe (in an admittedly weak sense of the term), of a tachistoscope, that it is making noise, yet not believe any proposition about it: she *attributes* noise-making to it, yet does not conceptualize it in the way required for

having a propositional belief about it, the kind of belief expressed in a complete declarative sentence such as 'The thing on the table is making noise'. She would then have no propositional belief about the instrument, the kind of belief that should unqualifiedly be called true (or false), such as that the tachistoscope is making noise. On this approach, what I am calling objectual belief is (or often is) better called *property attribution*. It is an attribution *to* the thing in question because of the kind of causal role that thing plays in grounding the attribution; and if it is not strictly speaking a belief, it does imply a disposition to form one, such as that the thing on the table is making noise.

5 Specifically, these are *doxastic* attitudes (from the Greek *doxa*, for 'belief'). A fear can be propositional and thereby cognitive, but it need not entail believing the proposition one fears is so, for example that the man approaching one will attack. Some might consider objectual awareness, say awareness of perfect symmetry, cognitive, at least when the person has the concept of relevant property. By contrast, desires, the paradigm *conative* attitudes, should not, I think, be taken to have propositional objects (e.g. 'to swim' in 'my desire to swim' does not express a truth or falsehood).

6 Perceptions that embody beliefs in the ways illustrated are also called *epistemic*, since the embedded belief is commonly considered to constitute knowledge. Their connection with knowledge is pursued in this chapter and others.

7 The distinction between simple and propositional perceiving and other distinctions drawn in this chapter are not always observed. At one point W.V. Quine says:

> think of "*x* perceives *y*" rather in the image of "*x* perceives that *p*". We say "Tom perceives the bowl" because in emphasizing Tom's situation we fancy ourselves volunteering the observation sentence "Bowl" rather than "Surface of a bowl," "Front half of a bowl," "Bowl and background," and so on. When we ask "What did he perceive?" we are content with an answer of the form "He perceived that *p*".
>
> (*Pursuit of Truth*, revised edn [Cambridge, MA: Harvard University Press, 1992], p. 65)

Notice that because seeing that (say) there is a bowl in front of one obviously entails seeing a bowl, it is no surprise that we are content with a report of the propositional perception even if we wanted to know only what object was seen: we get what we sought and more. It does not follow that simple seeing *is* or even entails propositional seeing. It is also worth noting that Quine is apparently thinking only of seeing here; for the other four senses, there is less plausibility in maintaining what he does.

8 The adage could not be taken to refer to simple seeing, for what we simply

see, say a glass or leaf or field, is not the sort of thing that can be believed (to be true or false). To be sure, seeing something, especially something as striking as golf ball-size hail, produces a *disposition to believe* certain propositions, say that this is a dangerous storm. But, by what seems an economy of nature, there are many things we are disposed to believe but do not. I have defended these points in detail in 'Dispositional Beliefs and Dispositions to Believe', *Noûs* 28 (1994), 419–34.

9 This applies even to full-scale philosophical dictionaries written by teams of experts, though such a work can provide concise statements of much valuable information. See, for example, the entries on blind sight and perception in Robert Audi (ed.), *The Cambridge Dictionary of Philosophy* (Cambridge: Cambridge University Press, 1995, 1999).

10 In the light of what has been said in this chapter so far we can accommodate much of what is plausible in the common view that, as D.M. Armstrong puts it:

> [perception] is an acquiring of knowledge or belief about our physical environment (including our own body). It is a flow of information. In some cases it may be something less than the acquiring of knowledge or belief, as in the cases where perceptions are entirely discounted or where their content has been confidently anticipated.
>
> (*Belief, Truth and Knowledge* [Cambridge: Cambridge University Press, 1973], p. 22)

First, I can agree that perception entails acquisition of *information*; the point is that *not all our information is possessed as the content of a belief*. Second, Armstrong himself notes an important way in which perception might fail to produce belief: it is "discounted," as, for example, where one is sure one is hallucinating and so resolutely refuses to accept any of the relevant propositions.

11 This is the kind of view developed in detail by Fred Dretske. See esp. *Knowledge and the Flow of Information* (Cambridge, MA: MIT Press, 1981).

12 The notion of normality here is not statistical; it implies that what is not normal calls for explanation. In the world as we know it, exceptions to the normality generalizations I propose seem at least rare; but the point is not that statistical one, but to bring out that the very concepts in question, such as those of seeing and knowing, have a connection in virtue of which explanation is called for if what is normally the case does not occur.

13 A property that something is seen as having need not be a property it actually has; but here seeing *as* is phenomenal, not doxastic. Roughly, the perceptual content represents what the object is like if it in fact has the properties it is seen as having.

14 A detailed discussion of the representationality of perception and the kind of content it has is provided by Fred Dretske in *Naturalizing the Mind* (Cambridge, MA: MIT Press, 1997). He deals with the sense in which perceptual content is external. If, loosely speaking, we call the perceived object the *objectual content* then simple perception obviously has a kind of external content; but as the object is "in" the experience, it might be considered a kind of content, as indeed it may for propositional and objectual perceptions as well. With this idea in mind, it is clear how the perceptually believed propositions themselves may also be conceived as having external content. I have discussed internal and external content in relation to such examples in 'Internalism and Externalism in Epistemology and Semantics', in Mark Timmons, John Greco, and Alfred R. Mele (eds.), *Rationality and the Good: Critical Essays on the Ethics and Epistemology of Robert Audi* (Oxford: Oxford University Press, 2007). (This responds to a challenge from Timothy Williamson, 'On Being Justified in One's Head', ibid., 106–122)

15 In speaking of justification that prevails, and of overall justification, I have in mind the kind appropriate to a rational person's believing the proposition in question, construed as roughly the kind such that when we believe a true proposition with that kind of justification then (apart from the kinds of case discussed in Chapter 10 that show how justified true beliefs *need not* constitute knowledge) we know it.

16 There are complexities I cannot go into, such as how one's competence figures. I am imagining here someone competent to tell whether a note is flat (hence someone not tone deaf): in general, if we are not competent to tell whether a kind of thing has a property or not, an experience in which it seems to have it may not justify us in believing it does. There is also the question of *what* the belief is about when the "object" is hallucinatory, a problem discussed shortly. Still other problems raised by this justification principle are discussed in Chapter 11 in connection with the controversy between internalism and externalism.

17 If, as is arguable, seeing that it is blue entails knowing that it is, then he does *not* see that it is, though he sees its blue color. But this entailment claim is far from self-evident. Suppose he clearly sees a blue bird and believes it is blue, but does not know that it is because of his frequent hallucinations. A moment before, he hallucinated such a bird; a moment later, he will again; and he realizes his senses have been playing such tricks on him. Still, he cannot help believing this bird is blue and believes that *on the basis of* clearly seeing it and its color in normal light. Might we say that he sees that the bird is blue, but does not know this? We cannot say that he "can't believe his own eyes," because he does; but if, in the normal way, they show him the truth and he thereby believes it, might he not see it through them?

2 Theories of perception

Sense experience, appearances, and reality

- **Some commonsense views of perception**
  Perception as a causal relation and its four main elements
  Illusion and hallucination

- **The theory of appearing**

- **Sense-datum theories of perception**
  The argument from hallucination
  Sense-datum theory as an indirect, representative realism
  Appraisal of the sense-datum approach

- **Adverbial theories of perception**

- **Adverbial and sense-datum theories of sensory experience**

- **Phenomenalism**
  A sense-datum version of phenomenalism
  Adverbial phenomenalism
  Appraisal of phenomenalism

- **Perception and the senses**
  Indirect seeing and delayed perception
  Sight and light
  Vision and the eyes

2 Theories of perception

Sense experience, appearances, and reality

Much has now been said about perception and its relation to belief, justification, and knowledge. What we have seen puts us in a position to frame a theory of what perception is, but we have not yet stated any such theory. Stating and assessing theories of perception is the task of this chapter.

Some commonsense views of perception

One natural thing to say about what it is for us to see the green field is appealingly brief. We simply see it, in an ordinary way: it is near and squarely before us; we need no light to penetrate a haze or a telescope to magnify our view. We simply see the field, and it may normally be taken to be pretty much as it appears. This sort of view, called *naive realism*, has been thought to represent common sense: it says roughly that perception is simply a matter of the senses telling us about real things.

The view is naive because it ignores problems of a kind to be described in a moment; it is a form of realism because it takes the objects of perception to be real things external to the perceiver, the sorts of things that are "out there" to be seen whether anyone sees them or not.

A more thoughtful commonsense view retains the realism without the naivety. It is quite commonsensical, for instance, to say that I see the field *because* it is before my open eyes and stimulates my vision, thereby *appearing* to me as a green, rectangular shape. Stimulating my vision is a causal relation. For instance, the field, by reflecting light, causes me to have the visual experience that is part of my seeing that very field. Moreover, the field apparently must cause my visual experience if I am to *see* it. As the more thoughtful commonsense view specifies that the object of perception must be a real external thing, we might call it a *perceptual realism*. Most—but not all—theories of perception incorporate perceptual realism.

To understand why perception must have a causal element, suppose I am looking at the field and, without my noticing, someone instantaneously drops a perfect picture of the field before my eyes. My visual experience might not change. What appears to me might look just as the field did. Yet I no longer *see* the field. Instead, I see a picture of it. (I do see the field *in* the picture, but

that is secondary seeing and not the kind we are talking about.) The reason I do not now see the field is roughly that it has no (causal) effect on my visual experience.

Perception as a causal relation and its four main elements

Examples like this suggest that *perception is a kind of causal relation* between whatever is perceived and its perceiver, wherein the object perceived produces a sensory experience in the perceiver. This is a plausible, commonsensical, and important point, though it does not tell us precisely what perception is. I call any theory of perception which incorporates the point *a causal theory of perception*. Most theories of perception are causal.

We can now better understand the four elements I have described as among those crucial in perception: the perceiver, the object perceived, the sensory experience in which the object appears to the perceiver, and the causal relation between the object and the perceiver, the relation wherein the object produces that experience. If you see the field, there is a distinctive way, presumably involving light transmission to your eyes, in which the field produces the visual sensory experience of a green, rectangular shape characteristic of your seeing it. If a picture of the field produces an exactly similar visual experience in the same way, it is the picture you see, not the field. Similarly, if you hear a piano piece, there is a special way in which it causes you to have the auditory sensations of chords and melody and harmony that go with it.

It is difficult, though fortunately not necessary for a general understanding of perception, to specify precisely what these causal paths from the object to the perceiver are. Some of the details are the business of the psychology and neurophysiology of perception. Others are determinable by philosophical inquiry. Philosophical reflection shows us, for instance, that not just any causal chain is the right sort for perception. Consider what is sometime called a *wayward* (*or deviant*) *causal chain*. Suppose the piano sounds cause a special machine, created by a prankster, to produce in me both temporary deafness and a faithful auditory hallucination of the piece. Then I do not *hear* it, though my sensory experience, the auditory experience I enjoy in my own consciousness, is just what it would be if I did hear it. Nor do I hear it if, though the sound waves reach my brain and cause me to believe a piano is playing just the piece in question, I have no auditory experience. Even such a highly informed inner silence is not musical.

Illusion and hallucination

We can make progress by pursuing the question of why naive realism is naive. Suppose there is a gray haze that makes the green field look gray. Or suppose the mouth of the cup I am holding appears, from a certain angle, as if it were an ellipse rather than a circle, or feels warm only because my hand is cold.

These are *perceptual illusions*, roughly perceptual experiences that (in the way illustrated) misrepresent the object of perception. They illustrate that things are not always as they seem. The cup is round and at room temperature.

Now imagine that the field burns up. I sorely miss its rich green and the spruce and maple, and on waking from a slumber in my chair I have a *hallucination,* roughly a sense experience qualitatively like a perceptual one but not of an external object. In this case, my (hallucinatory) visual experience is just as it would be if I were seeing the field. Here the grass I seem to see is not there at all. The point is not that something I see is not as it seems (as in the case of illusion) but that there seems to be something where there is nothing. With illusion, as illustrated by a partly submerged stick's looking bent, what is there is perceived distortedly; with hallucination, it appears that something is perceived when nothing is. Illusions and hallucinations are possible for the other senses too. When they occur, we do not just see (or hear, taste, smell, or touch) the object. Either we do not see it as it is or (perhaps) we do not see anything at all. Not everything we perceive is as it appears to be, and naive realism does not explain why.

One way to deal with illusion and hallucination is to stress how they show the need to distinguish appearance from reality. In a visual illusion, one sees something, but it does not appear as it really is, say circular. In a hallucination, nothing need appear to me at all, and if anything does, it is in reality even less what it appears to be than is the object of an illusion, or is not what it appears to be at all: instead of a blue spruce tree's appearing blue to me, for instance, it is as if the conical section of space where it stood appears "bespruced."

The theory of appearing

The sort of account of perception just sketched as an improvement over naive realism has been called *the theory of appearing*: it says roughly that perceiving an object, such as a book, is simply its appearing to one in a certain (sensory) way. It may, for instance, appear to be rectangular. Thus, one perceives it—in this case, sees it—*as* rectangular. The theory can also provide the basis of an account of the sort of experience we have in hallucination as opposed to normal perception: that experience, too, the theory takes to be a case of something's appearing to one to have a set of properties; the object that appears is simply a different kind.[1]

The theory of appearing is initially plausible. It incorporates much reflective common sense, for instance the view that if one sees something, then it appears to one in some way, say as a red barn or at least as a visually experienced rectangular patch. The theory says nothing, however, about the need for a causal relation between the object and its perceiver (though it allows that there be one). If, consistently with its commonsense motivation, one stipulated that the crucial relation of appearing to the perceiver to have a

property—say to be rectangular—is a causal relation, one would then have a different theory (of a kind to be discussed shortly).

In addition to the question of how the theory can do justice to the causal element in perception, it faces a problem in accounting for hallucinations in which there is no object to appear to the person at all. I could, after all, hallucinate a green field when I see nothing, say because it is pitch dark or my eyes are closed. In such an *empty hallucination*—one that occurs despite my perceiving nothing—what is it that appears green to me? One plausible answer is given by a quite different theory of perception. Let us explore that view.

Sense-datum theories of perception

Once we think seriously about illusion and hallucination, we may begin to question not only naive realism but also any kind of *direct realism* in the theory of perception, any perceptual realism which, like the theory of appearing, says that we see (or otherwise perceive) external objects directly, rather than *through* seeing (or at least visually experiencing) something else. After all, not only do light rays come between us and what we see, there are also brain events crucial for seeing. Perhaps these events or other intermediaries in perception produce or indicate an interior object, presumably a mental object that plays an intermediary role in perception.

Hallucination illustrates most readily how such an intermediary might figure in understanding perception. Imagine that you vividly hallucinate the field just as it would be if it were before you. This seems quite possible. If such a "faithful" hallucination occurs, your visual experience—roughly, what you are aware of in your visual consciousness—is exactly like the experience you have when you see the field. Does it not then seem that the difference between ordinary seeing and visual hallucination is simply in what *causes* the visual experience, rather than in what you directly see? When I see the field, *it* causes my visual experience. When I hallucinate it, something else (such as my deep desire to have it back) causes my visual experience. But apparently what I directly see—the immediate object of my visual experience—is the same object in both cases. This point presumably explains why my visual experience is qualitatively the same whether I am hallucinating the field or really seeing it. If it were not the same, we could not say things like 'It was exactly as if I were seeing the tree in normal light'.

The argument from hallucination

We might develop these ideas by considering an argument from hallucination. It consists of two connected arguments. The first constituent argument attempts to show a parallel between hallucination and ordinary perception:

1 A perfectly faithful (visual) hallucination of a field is qualitatively just like, and so intrinsically indistinguishable from, an ordinary experience of seeing that field, that is, not distinguishable from it just in itself *as* a visual experience, as opposed to being distinguishable through verifying one's visual impression by touching the things around one.[2]

Hence:

2 What is *directly* "seen," the immediate object of one's visual experience, is the same sort of (non-physical) thing in a perfect hallucination of a field as in an ordinary experience of seeing a field.

But—and we now come to the second constituent argument, which builds on (2) as its first premise—clearly:

3 What is directly seen in a hallucination of a field is not a field (or any other physical thing).

Indeed, no field is seen at all in a hallucinatory visual experience, so (3) seems plainly true. Hence, putting (1)–(3) together, we may infer that:

4 What is directly seen in an ordinary experience of seeing a field is not a field.

The overall idea is that when we ordinarily see an everyday perceptible object such as a field, we see it through seeing—in the sense of visually experiencing—something else *directly*: something not seen *by* seeing anything else. What we see directly—call it a *sense-datum*—might be an image. One may prefer (as some philosophers do) to say that we do not in any sense *see* such things but are only visually acquainted with them. To simplify, let us bear this alternative in mind but use the more natural term 'see'.

Just what is directly seen when one sees a field, then, and how is the field *in*directly seen? Why not say that what is *directly* seen is a two-dimensional object consisting of the same sorts of colors and shapes one sees in the hallucinatory experience? After all, nothing, not even (physical) light, intervenes between us and them. There is no "space" for intermediaries. Hence, no intermediaries can misrepresent these special objects. These objects are apparently internal to us: as traditionally conceived, they could exist even if—as Descartes, in his *Meditation I,* supposed to be possible—we were disembodied minds in an otherwise empty world. The only space they need is in the mind. Yet we do see the field *by* seeing them; hence, we see it indirectly, at least in the sense that we see it by having a visual experience that is directly of something else.

The idea that experiencing sense-data is required for perception is nicely expressed in Emily Dickinson's poem 'I Heard a Fly Buzz When I Died'. In the final moment of her terminal experience:

There interposed a fly,

With blue, uncertain stumbling buzz,
Between the light and me;
And then the windows failed, and then
I could not see to see.

The external light from the window blocks her eyesight, but this leaves inner seeing—portrayed here as a necessary condition for ordinary seeing—still possible. Until the end, she can see *to* see. It is sense-data that are conceived as the direct objects of such inner sight.

A sense-datum theory is perfectly consistent with a causal theory of perception: the field produces the colors and shapes in my visual consciousness in a way that fully accords with the view that perception is a causal relation between something external and the perceiver. Perception is simply a *mediated*, hence *indirect*, causal relation between external objects we perceive and us: the object produces the mediating colors and shapes that appear in our visual fields, and, through seeing them, we see it.

The theory I am describing is a version of a *sense-datum theory of perception*. Such theories are so called because they account for perception by appeal to a view of what is directly given *in* sense experience, hence is a *datum*, a given, for such experience—the sort of thing one is visually aware of in hallucinating a field. This sense-datum thesis (unlike the phenomenalist sense-datum view to be discussed shortly) is a realist view; but its realism, by contrast with that of naive realism and the theory of appearing, is indirect.[3]

Sense-datum theory as an indirect, representative realism

A sense-datum theory is a kind of *representative realism* because it conceives perception as a relation in which sense-data represent perceived external (hence real) objects to us. On some conceptions of sense-data, they are copies of those objects: shape for shape, color for color, sound for sound. John Locke held a view of this kind (and in 1689 published it in *An Essay Concerning Human Understanding*, especially Books II and IV), though for him sense-data are copies ("resemblances") only of the *primary qualities* of physical things—solidity, extension (in space), shape, and mobility—not of their *secondary qualities*, above all colors, sounds, smells, and tastes. (He took the primary qualities to be objective and of the kind that concern physical science; and he considered the secondary ones to be in a sense subjective, not belonging to physical things but something like representational mental elements that they cause in us. Color, then, disappears in the dark, though the physical object causing us to see it is not changed by the absence of light.) Our question is whether any sense-datum version of representationalism is sound, and we need not pursue the interesting question of how these two kinds of qualities differ.

Sense-datum theories have had illustrious defenders down to the present

age. The theory has also had powerful opponents. To appreciate it better, let us first consider how it takes perception to be indirect. Sense-datum theorists might offer several reasons to explain why we do not ordinarily notice the indirectness of perception (I speak generally, not solely of Locke's version of the theory). Here are two important reasons.

First, normally what we directly see, say colors and shapes, roughly corresponds to the physical objects we indirectly see by means of what we see directly. It is only when there is an illusion or hallucination that we are forced to notice a discrepancy between what we directly see and the object commonly said to be seen, such as a book.

Second, the beliefs we form on the basis of perception are formed spontaneously, not through any process requiring us to consider sense-data. Above all, we do not normally *infer* what we believe about external objects we see from what we believe about the colors and shapes we directly see. This is why it is easy to think we "just see" things, directly. Perceiving is not inferential, and for that reason (perhaps among others) it is not *epistemically indirect*, in the sense that *knowledge* of external objects or beliefs about them are indirect, in the sense that they are based on knowledge of sense-data, or beliefs about them. On a plausible sense-datum view, I know that the field is green through *having* rectangular green sense-data, not through *inference from* propositions about them.[4]

It is apparently true that, as a sense-datum view may allow, perception is not inferential or epistemically indirect in the way inferentiality would imply. But, for sense-datum theorists, perception is nonetheless causally and objectually indirect. The perceived object is presented to us via another object, though not by way of a *premise*. These theories are *causally indirect*, then, because they take perceived physical objects to cause sensory experience, say of colors and shapes, *by* causing the occurrence of sense-data, with which we are directly (and presumably non-causally) acquainted in perceptual experience. Perception is also *objectually indirect* because we perceive external things, such as fields, *through* our acquaintance with other objects, namely sense-data. Roughly, we perceive external things through perceptual acquaintance with internal things.

Despite the indirectness of perception in these two respects, a sense-datum theorist need not deny that we normally do not use information about sense-data to arrive at perceptual beliefs inferentially, say by an inference from my directly seeing a grassy, green rectangular expanse to the conclusion that a green field is before me. Ordinarily, when I look around, I form beliefs about the external environment and none at all about my sensory experience. That experience causes my perceptual beliefs, but what they are *about* is the external things I perceive. It is when the colors and shapes do not correspond to the external object, as when a circle appears elliptical, that it seems we can understand our experience only if we suppose that the direct objects of sensory experience are internal and need not match their external, indirect objects. Is the sense-datum view, however, correct on this point?

Appraisal of the sense-datum approach

Let us focus first of all on the argument from hallucination, whose conclusion suggests that what is directly seen in visual perception of external objects are sense-data. Suppose I do have a hallucination that is qualitatively just like, and intrinsically indistinguishable from, the normal experience of seeing a field. Does it follow that what is directly seen in the hallucination is the same sort of thing as what is directly seen in the normal experience? At least two problems confront the sense-datum theory here.

First, why must *anything* be seen in a hallucination? Imagine that I hallucinate the burned-up field. I might get up, still half asleep, and, pointing to the area, cry out, "It's regrown!" You might conclude that I *think* I see the field again. My *initial* reaction to realizing I had hallucinated the field might be that, hallucination or no, I *saw* it. But I might just as easily slump back in my chair and mumble that I wish I had seen it.

A compromise view would be that I saw the hallucinated grass *in my mind's eye*. But suppose I did see it in my mind's eye, and again suppose that the hallucination is intrinsically just like the ordinary seeing. Does it follow that what I directly see in the ordinary experience is the same as what I see in the hallucination, namely, something in my mind's eye? It does not. The notion of seeing in one's mind's eye is metaphorical, and such seeing need not imply that there is any real object seen, in or outside the mind. However vividly I may, in my mind's eye, see myself standing atop a giant pyramid in Toronto, that city has no pyramid, nor need there be any pyramidal object in my mind.

There is a second reason to resist concluding that something must be directly seen in hallucinations. Recall that my seeing a green field is apparently a causal relation between a sensory experience in me and the field that produces the experience. If so, why should the possibility that a hallucination can mimic my seeing the field tell us anything about what is directly seen (or is an object of visual acquaintance) when one sees that field? It is not as if we had to assume that only an experienced *object* can produce, or causally sustain, the relevant sensory experience, and must then conclude that it is an internal perceptual object, since there is no other candidate. Many things can have more than one causal basis, and the sense-datum theorist has no argument to show that only an internal perceptual *object*, or an acquaintance with it, as opposed, say, to an abnormality in the visual cortex (which need not be an object at all), can cause or sustain the hallucinatory experience.

Moreover, from the similarity of the internal, experiential elements in the hallucination and the internal ones occurring in genuine perception, one might as well conclude that since the ordinary experience is one of seeing only an external rather than an internal object, the hallucinatory experience is different only in the absence of the external object. Rather than add to the components that seem needed to account for the ordinary experience, we

subtract one that seems needed to account for the hallucination. This yields a more economical theory of perception.

An analogy may help. Compare trying to infer facts about how we see an original painting from facts about how we see it in a photo of it. From the indirectness of the latter seeing, it certainly does not follow that ordinary seeing of the painting is indirect. And even if a photographic viewing can be so realistic that it perfectly mimics an ordinary viewing, it does not follow that photographic, two-dimensional objects are components in ordinary seeing. Similarly, no matter how much like ordinary experiences hallucinations can seem, it does not follow that the former have all the internal elements (roughly, mental or mind-dependent elements) of the latter.

It will help to consider a different analogy. Two perfect ball-bearings, by virtue of identical diameter and constitution, can be qualitatively alike and so intrinsically indistinguishable, yet they can still differ, one being on your left and one on your right. Their intrinsic (i.e., non-relational) properties can thus be identical, while their *relations* (to you) differ. Hence they *do* differ in their relational properties. Similarly, the hallucination of a field and the ordinary visual experience of a field can be qualitatively alike, and so intrinsically indistinguishable, yet differ in their relations to me or to other things. One of them, the visual experience of a field, may be an element in a perceptual relation to the field; and the experience we call hallucination, which is not based on perceiving the external object hallucinated, may *not* be an element in any perceptual relation to the field, but only something I undergo (an element simply "in" me, on the plausible assumption that it is mental).

To account for the difference between the two kinds of experience, we might say this: the visual experience indicates the presence of an external thing; the hallucinatory experience, though intrinsically just like the visual one, only pretends to indicate an external thing to me. Thus, for all the argument from hallucination shows, the ordinary experience of seeing might be a relation to an object such as a green field, namely the relation of directly seeing, while the hallucinatory experience of a green field is not a relation to that field, such as being an internal copy of it, or even a relation to any other object, such as a perceiver.

The points just made about the argument from hallucination indicate that it is not sound. Its first premise, (1), does not entail (2) the conclusion drawn from it. Nonetheless, the argument poses serious problems for alternative theories. What explanatory account of hallucinations and illusions besides the sense-datum account might we adopt? To see some of our alternatives, it is best to begin with illusion rather than hallucination.

Recall the mouth of the cup viewed from an angle. A sense-datum theory will say we directly see (or anyway experience) an elliptical shape and indirectly see the cup. The theory of appearing, however, can also explain this: it reminds us that things need not be what they appear to be and says simply that they can appear elliptical even if they are round.

Adverbial theories of perception

One could also combine the causal element in the sense-datum approach with the direct realism of the theory of appearing and move to a third theory, one that says the cup causes us to see it directly, rather than through producing sense-data in us, yet (because of our angle of vision) we see it as if it were elliptical. To avoid suggesting that anything in one's experience need *be* elliptical, one could take this to mean that the cup visually "appears elliptically" to us. Here the adverb 'elliptically' modifies the verb 'appears' and describes a *way* in which we visually experience the cup. It does not imply that there is an object that appears to us and *is* elliptical.[5] Let us explore this idea in relation to the theory associated with it.

It should now be clear why we need not grant (what sense-datum theorists sometimes seem to assume about perception) that in order for an object to appear a given way to us there must *be* something we see that *is* that way, for instance an elliptical sense-datum. Suppose that one says simply that the cup appears elliptically, using this adverb to designate, from the perceiver's point of view, *how* one visually experiences it: elliptically. To say it appears elliptically is roughly to say it appears in the way an ellipse does, viewed from directly above its center, as opposed to the way a circle does when so viewed.

If this adverbial interpretation of such statements as 'I see an ellipse' seems artificial, consider an ordinary analogy. If I say I have a fever, no one could plausibly insist that there is an object, a fever, which I have. 'I have a fever' is a way of saying I am feverish, that is, my body is above a certain temperature. What our language seems to treat as a statement of a relation to an object, a fever, is really an ascription of a property: the property of having a temperature above a certain level. Just as 'having a fever' can ascribe a certain temperature, 'seeing elliptically' (in illusional and hallucinatory cases) can ascribe a certain visual experience.

On the basis of this move, one can construct what is called *the adverbial theory of perception*. Unlike the theory of appearing, which takes perception to be an unanalyzable relation in which things appear to us as having one or more properties, an adverbial theory conceives perception as an analyzable way of experiencing things. In what may be its most plausible form, it says roughly that *to perceive an object is for that object (in a certain way) to produce in one a sensory experience of it*: to cause one's experiencing it in a certain qualitative way, say to see a stick as straight (or, given the illusion induced by partial submersion, as bent). Both theories are, however, direct realist views. In both, the perceptual object appears to the perceiver without an intermediary. Other similarities (and some differences) between the two theories will soon be apparent.[6]

The adverbial theorist stresses that we see (or otherwise perceive) things in a particular qualitative way and that they thus appear to us in that way. Often they appear as they are; sometimes they do not. In each case they

are seen directly, not through intermediaries. Even if I do not see the cup as circular, I do see *it*: it is seen directly, yet appears elliptically.

So far, so good, perhaps. But what about hallucinations? Here the adverbial theory again differs from the theory of appearing. Unlike the latter, it denies that all sensory experience is *of* some object. The importance of this denial is not immediately apparent, perhaps because we suppose that usually a person visually hallucinating does see *something*. Consider Shakespeare's Macbeth, distraught by his murder of Duncan, hallucinating a dagger that seems to him to hover in mid-air:

> Is this a dagger which I see before me,
> The handle toward my hand? Come, let me clutch thee.
> I have thee not, and yet I see thee still.
> Art thou not, fatal vision, sensible
> To feeling as to sight? or art thou but
> A dagger of the mind, a false creation,
> Proceeding from the heat-oppressed brain?
> I see thee yet, in a form as palpable
> As this [sword] which now I draw.
>
> (Act II, scene i)

Perhaps Macbeth does see something in the place in question, say the wall behind "the dagger" or at least a chunk of space where it hovers. Thus, to explain what it is for Macbeth to experience "daggerly," an adverbial theorist might posit an "object" where the "dagger" seems located, if only the section of space where it seems to be. On one view, this object might be thought to play a role in causing him to have daggerish visual sensations. But there need not be any such object in a hallucination. For the theory of appearing, too, the space before him, despite being transparent, might somehow appear to him to be a dagger.

Supposing we accept this adverbialist account, what happens if it is pitch dark and Macbeth's hallucination is therefore *empty*, in the sense that there is nothing he sees, and hence nothing to serve as an object visually distorted into an apparent dagger? Then, whereas the theory of appearing may have to posit a sense-datum (or other special kind of object) to serve as what appears to be a dagger, the adverbial theory can deny that there is *any* kind of object appearing to Macbeth. It may posit some quite different account of his "be-daggered" visual experience, such as a psychological account appealing to the influence of drugs or of his "heat-oppressed brain."

Is it really plausible to hold, with the adverbial theory, that Macbeth saw nothing at all? Can we really explain how the normal and hallucinatory experiences are intrinsically alike without assuming they have the same direct objects? In the light of the special case of empty hallucination, the sense-datum theory may seem the most plausible of the three. It provides an object of Macbeth's visual experience even if that occurs in utter darkness,

whereas the adverbial theory posits no perceptual objects at all in empty hallucinations. Moreover, the sense-datum view postulates the same sort of direct object for ordinary perception, illusion, and hallucination, whereas the theory of appearing does not offer a uniform account of their direct objects and must explain why entities such as sense-data do not occur in normal perception as well as in empty hallucination.

Perhaps, however, the hallucination problem seems more threatening than it should to the adverbial theory because hallucinations are felt to be *perceptual* experiences and hence expected to be *of* some object. But as we have seen, although hallucinatory experiences can be intrinsically indistinguishable from perceptual ones, all that can be assumed is that the former are *sensory experiences*. Hallucinatory experiences, on the adverbial view, are simply not cases of perceiving, at least not in a sense requiring that any object appear to one.

Thus, nothing at all need appear to one in hallucinations, though it may *appear to the subject* that something is there. The hallucinator may then be described as having a visual sensory experience, but—as nothing is perceived—not a genuine perceptual experience.

Adverbial and sense-datum theories of sensory experience

A perceptual experience is always sensory, and normally a sensory experience of the sort we have in perceiving is genuinely perceptual. But, as hallucination shows, a kind of short-circuit can cause the sense-receptors to produce sensory experience that is not a normal perceptual experience (or even part of one). It is important to consider the debate between adverbial and sense-datum theories in relation to sensory experience. Both theories take such experience to be essential to perception, and both offer accounts of sensory experience as well as of perception.

The most natural thing for adverbial theorists to say about hallucinatory experience is that it is not genuinely perceptual, but only sensory. They might, however, say instead that when a perceptual experience is hallucinatory, it is not a case of *seeing* (except perhaps in the mind's eye, or perhaps in the sense that it is seeing colors and shapes conceived abstractly as properties and not as belonging to sense-datum objects). The former description accords better with how seeing is normally understood: normally, we cannot be said to see what is not there.

The theory suggested by these responses to the hallucination problem might be called *the adverbial theory of sensory experience*. It says that having a sensory experience, such as a hallucination of a green field, is experiencing in a certain *way*, for example visually experiencing "green-fieldly." Our commonsense assumption is that hallucination is not usual (for normal people) and that most of our vivid sensory experiences are genuinely perceptual. They are of, and thus caused by, the external object(s) apparently perceived.

But some sensory experiences are neither genuinely perceptual nor externally caused. People having them are in, say, a vision-like state, and what occurs in their visual cortex may be the same sort of process that occurs when they see things. Yet they are not seeing, and their visual experience typically has an internal cause, such as an abnormal emotion.

May we, then, regard sense-datum theories of perception as refuted by the points just made in criticism of the argument from hallucination and on behalf of the suggested adverbial theory and the theory of appearing? Certainly not. We have at most seen how one major argument for a sense-datum theory of perception fails and how alternative theories of perception can account for the apparently central elements in perception: the perceiver, the (ordinary) object perceived, the sensory experience, and the causal relation between the second and third.

Indeed, supposing that the argument from hallucination fails to show that sense-data are elements in normal everyday perception, sense-data might still be needed to account for non-perceptual sensory experience (sometimes loosely called perceptual experience because it is characteristic of perception). In this limited role, one might posit a *sense-datum theory of non-perceptual sensory experience*: such experience is considered direct acquaintance with sense-data.

A sense-datum view may seem preferable to an adverbial theory of sensory experience. For one thing, there is something unsatisfying about the idea that even in a visual hallucination so vivid that, if one did not suspect error, one would stake one's life on the presence of the hallucinated object, one sees nothing, except either in one's mind's eye, or in a sense of 'see' which does not require that any object be seen. Still, perhaps there is such a sense of 'see'.

There is another aspect of the controversy. It concerns the *metaphysics* associated with adverbial and sense-datum theories of any kind, specifically the sorts of things they require us to take as fundamental realities. In this respect, the adverbial theories of perception and sensory experience have an advantage over the counterpart sense-datum theories: the former do not posit a *kind* of object we would not otherwise have to regard as real. From the adverbial perspective, the objects that perception and sensory experience involve are simply perceivers and what they perceive. These are quite familiar entities which we must recognize and deal with anyway.

Sense-data are quite different from ordinary (presumably physical) objects of perception. Sense-data are either mental or at least depend for their existence on the mind of the subject. Yet they are unlike some mental phenomena in that no plausible case can be made for their being really brain phenomena, since they have properties, for instance being green and perfectly rectangular, not normally found in the brain.[7]

Moreover, there are difficulties in fully understanding sense-data in any terms. Is there, for instance, even a reasonable way of counting them? Suppose my image of the green field gradually gets greener. Is this a sense-datum changing or a new one replacing an old one? There seems to be no

way to tell. If there is no way to tell, how can we ever be sure we learn more about a sense-datum than what initially appears to us in experiencing it: how can one distinguish learning something more about *it* from learning about something new?[8]

Problems like these also affect the theory of appearing insofar as it must posit sense-data or similar entities to account for hallucinations. To be sure, such problems can also beset our understanding of ordinary objects. Can we always distinguish a mountain with two peaks from two mountains, or one snarled barberry bush from two with intertwined roots? But these problems seem less serious, if only because there is no question that there are *some* things of the physical kind in question. The corresponding problems may in the end be soluble for sense-data, but they at least give us some reason to prefer a theory that does not force us to regard sense-data as the only objects, or as even among the objects, we are directly aware of when we see, hear, touch, taste, and smell.

Phenomenalism

If some philosophers have thought that perception can be understood without appeal to sense-data, others have conceived it as understandable in terms of sense-data alone as its objects. This view has the advantage of being, in at least one way, simpler than the adverbial and sense-datum theories. But the view is motivated by other considerations as well.

A sense-datum version of phenomenalism

Think about the book you see. It is a perceptible object. Suppose we may conceive a real perceptible object as a perceptible object that is as it is, independently of what we think it to be. Still, real perceptibles, such as tables and chairs and books, are also plausibly conceived to be, by their very nature, *knowable*, at least in being experienceable. Indeed, it is doubtful that real objects of this sort could be unknowable, or even unknowable through the senses if lighting and other perceptual conditions are good. Now suppose we add to these ideas the assumption that our only genuine, certain knowledge of perceptibles is restricted to what directly appears to us and would be as it is even if we should be hallucinating. And what more does appear to us besides colors, shapes and other sensory (sensible) properties? Further, how do we know that this book, for example, could even exist without someone's perceiving its sensory properties? Certainly we cannot *observe* the book existing unperceived. If you observe it, you perceive it.

Moreover, if you imagine subtracting the book's sensory properties one by one—its color, shape, weight, and so on—what is left of it? This is not like peeling an apple, leaving its substance. It is like stripping layer after layer from an onion until nothing remains. Might we not conclude, then, that the book is not only *known by* its sensory properties, as the other theories of

perception also hold, but also *constituted by* a stable collection of them—by visual, tactual, and other sense-data that recur in our experience, confronting us each time we have the sense-data corresponding to, say, a certain bookcase in our home? Similarly, might it not be that to see the book *is* simply to be visually acquainted with such a stable collection of sense-data?

George Berkeley argued from a variety of angles that this is indeed what a perceptible object is. This view (which Berkeley developed in detail in his *Treatise Concerning the Principles of Human Knowledge*, published in 1710) is a version of what is often called *phenomenalism*, as it constructs external objects out of phenomena, which, in this use of the term, are equivalent to sense-data. The view is also considered a kind of *idealism*, since it construes physical objects as ideal, in the sense of being composed of "ideas" (an old term for sense-data) rather than material stuff that would exist even if there were no minds and no ideas.[9]

Adverbial phenomenalism

Phenomenalism as just described is focused on the nature of perceptible objects but implies a related view of perception. In the sense-datum version of phenomenalism we have been examining, the associated account of perception retains a sense-datum theory of sensory experience, but not a sense-datum theory of perception. The latter view posits external objects as causes of the sense-data experienced in ordinary perception, whereas sense-datum phenomenalism says that physical objects *are* collections of sense-data.

Using the adverbial theory of sensory experience, one might also formulate an *adverbial phenomenalism*, which constructs physical objects out of sensory experience alone and says that to see (for instance) a green field is to experience "green-fieldly" in a certain vivid and stable way. To see a thing is to have a visual experience that predictably occurs under certain conditions, say when one has the experiences of walking out on the porch and looking ahead.

On this phenomenalist view, perception can occur without even sense-data; it requires only perceivers and their properties. Sense-datum versions of phenomenalism, however, have been more often discussed by philosophers, and I will concentrate on them.

Whereas the sense-datum theory is an indirect realism, phenomenalism is a *direct irrealism*: it says that perceptual objects are directly perceived, but it denies that they are real in the sense that they are mind-independent and can exist apart from perceivers. This is not to say they are not perceptually real—real items *in sensory experience*. The point is that they are not metaphysically real: things that are "out there," which are the sorts of things we think of as such that they would exist even if there were no perceivers.

Phenomenalism does not, then, deny that physical objects exist in the sense that they are both stable elements of our experience and governed by

causal laws, such as those of physics. Nor does it deny that there can be hallucinations, as when certain experiences, like those presenting Macbeth's hallucinatory dagger, are too unstable to represent a physical object, or occur in only one mode, such as vision, when they should have tactile elements as well, such as a cool, smooth surface. What phenomenalism denies is that physical objects are real in the classical sense, implying that their existence is independent of our experience.

One naturally wonders why things would not go in and out of existence depending on whether they are experienced, and why, when they do exist, they obey the laws of physics, which certainly do not seem to depend on our minds. Berkeley did not neglect to consider what happens to things when we cease to perceive them, as when we leave a book in an empty room. His answer has been put in the following exchange:

> There was a young man who said "God
> Must think it exceedingly odd
> If he finds that this tree
> Continues to be
> When there's no one about in the quad."

Reply:

> Dear Sir:
>
> Your astonishment's odd:
> I am always about in the quad
> And that's why the tree
> Will continue to be,
> Since observed by, Yours faithfully, God.

If the very existence of external objects is sustained by divine perception, it is not difficult to see how their behavior could obey laws of nature that are divinely ordained.

A phenomenalist need not be a theist, however, to offer an account of the stability of external objects and their lawful behavior. John Stuart Mill, writing in the same epistemological tradition as Berkeley but without any appeal to God, considered external objects *permanent possibilities of sensation*. To say that the book is in the room when no one is there to perceive it is to say that there is a certain enduring possibility of the sensations, where having those sensations in a certain stable way constitutes perceiving such a book. If one enters the room and looks in the appropriate direction, that possibility should be realized. By contrast, if one merely hallucinated, there would be no reason to expect this. A phenomenalist can, however, be more radical and take objects not to have any kind of existence when unperceived. They are born and die with the experiences in which they appear.

Appraisal of phenomenalism

Unlike the sense-datum theory of perception, phenomenalism is only occasionally defended by contemporary philosophers. But it has had major influence. Moreover, compared with the sense-datum theory, it is more economical and in that way simpler. Instead of perceivers, sense-data, and external objects, it posits, as the things figuring in perception and sensory experience, just perceivers and sense-data. Indeed, adverbial phenomenalism does not even posit sense-data.

As a theory of perception, then, phenomenalism has fewer objects to analyze and interrelate than do the other theories we have discussed. In addition, it appears to bridge the most important gap between sensory experience and perception of objects: since the objects are internal and directly experienced, it seems natural to say that they must be as they appear to be—we see all there is of the surface facing us and in principle can see all there is to them as physical objects. On the other hand, for the external objects of common sense, whose reality is independent of perceivers, phenomenalism (if nontheistic) must substitute something like permanent possibilities of experience. Thus, the bare-bones appearance of the theory is illusory. Even that metaphor is misleading; for our bodies are also collections of sense-data; even the flesh itself is not too solid to melt into the sensations of its perceivers.

It is tempting to reject phenomenalism as preposterous. But if we do, we learn nothing from it. Let me pose just one objection from which we learn something important about the relation between sense experience and external objects. The theory says that a book, for instance, is—or at least that its presence is (necessarily) equivalent to—one's having or potentially having a suitably stable collection of sense-data, and that seeing it is being visually acquainted with them. If this is a correct analysis of what seeing a book is, then there is a combination of sense-data, sensory items such as colors and shapes in one's visual field, such that if, under appropriate conditions, these elements occur in me, then it follows that I see a book. But surely there is no such combination of sense-data—a point that is important for understanding skepticism. No matter how vividly and stably I (or anyone) may experience the colors and shapes appropriate to a book, it does not follow that anyone sees one. For it is still possible that I am just hallucinating one or seeing something else *as* a book.[10]

This kind of hallucination remains possible even if I have supporting tactual experiences, such as the smooth feel of paper. For even the sense of touch can be hallucinatorily stimulated. Thus, seeing a book is not *just* having appropriate booklike experiences, even if it is *partly* this, and even though, as phenomenalists may hold, there is no experienceable difference between a sufficiently stable combination of booklike sense-data and an independently real material book. Still, if seeing a book is not equivalent to any such collection of sensory experiences, phenomenalism fails to account for perception of ordinary objects. If there are objects for which it holds, they are not the kind we have in mind in seeking an account of perception.

Perception and the senses

I want to conclude this chapter by indicating some remaining problems about perception. I have already suggested that adverbial theories, sense-datum theories, and the theory of appearing provide plausible accounts of perception, though I consider some version of the first kind prima facie best and I leave open that some theory different from all of them may be better than any of them. I have also suggested that some perceptually grounded beliefs fail to be justified, and that, even when justified and true, they can fail to constitute knowledge. There are two further kinds of problems we should explore. One kind concerns observation, the other the relation of perception to the five senses.

Indirect seeing and delayed perception

Observing something in a mirror can count as seeing it. Indeed, it illustrates the sort of thing ordinarily considered seeing something indirectly, as opposed to seeing it by seeing sense-data. We can also speak of seeing through telescopes and other instruments of observation, again indirectly. But what if the object is microscopic and colorless, yet appears to us through our lens as gray? Perhaps we see it, but not quite as it is.

If we see a microscopic object at all, however, there must be some respect in which what we see it by is faithful to it or at least represents it by some relation of causal dependence—sometimes called *functional dependence.* This relation is perhaps more aptly termed a *discriminative dependence*, since perceptual experience seems to vary as a function of certain changes in the object, as where a bird's moving leftward is reflected in a movement of the image in our binoculars, yet in a systematic way that enables us to discriminate it from its environment. But what we see a thing *by*, such as color and shape, need not be faithful in all respects. A green field can look black at night; we are nonetheless seeing it. Moreover, we can see something move in the field even if its color *and* shape are distorted.

How much *correspondence* between an object and our sensory impressions representing it to us is required in order for us to see it (or hear it, touch it, and so on)? There may be no answer to this question that is both precise and highly general. The cases vary greatly, and many must be examined in their own terms.

Observation of faraway objects poses further problems. Consider seeing the nearest star. It is commonly taken to be about four light years away. Presumably we see it (if at all) only as it *was*. For the sense-datum theory, we have a sense-datum produced by it as it was. On the adverbial view, we are sensing "starly" in the way we would if we received the relevant visual stimuli at the time the star produced them. If, however, we see it only as it was, do we see *it*, or just its traces?

Suppose that unbeknownst to us the star exploded two years ago. Is it not odd to say we now see it, as opposed to seeing traces of it (as it was)? The

latter view is preferable, on the ground that if we unqualifiedly see something now, it exists now. But this point is compatible with the view that even though we may see a thing that exists now only *as* it was, we still literally see it now, just as, if we see a cup as elliptical when it is in fact round, we still do see the cup. In the case of the star, *if* the causal connection between it and our sensory experience representing it required no elapse of time, we would discriminate it sufficiently to see it as it now is.

Similar points hold for everyday seeing, since there is still some temporal gap, and for hearing. But if I can see the field only as it was a fraction of a second ago, can I still know that it is now green? I think so, provided there is no reason to believe its color has suddenly changed (but this is something to be reconsidered in the light of our discussion of skepticism in Chapter 13). The same is not clear for the star: may we know by sight alone that it exists now, when it would take about four years for us to realize that the light that had been emitted was its last? This seems doubtful, but it may depend on how likely it is that a star of this kind might have burned out during the period in question. If we knew that such stars last billions of years and that this one is only a few million years old, we might plausibly think we know it still exists. It is plain, however, that understanding perception and perceptual knowledge in these sorts of cases is not easy.

Sight and light

We normally regard seeing as intimately connected with light. But must seeing involve light? Suppose you could step into a pitch-dark room and have precisely the experiences you would have if it were fully lighted. The room would thus *look* to you just as it would if fully lighted, and you could find any unobscured object by looking around for it. Would this not show that you can see in the dark? If so, then the presence of light is not essential to seeing.

However, the case does not establish quite this much. For seeing is a causal relation, and for all I have said you are just vividly hallucinating precisely the right things rather than seeing them. But suppose you are not hallucinating and that if someone covered a coin you see with lead or covered your eyes, you would no longer have a visual experience of a coin. In this case, it could be that somehow the coin affects your eyes through a mechanism other than light transmission, yet requiring an unobstructed path between the object seen and your eyes. *Now* it begins to seem that you are seeing. You are responding visually to stimuli that causally affect your eyes. Yet their doing so does not depend on the presence of light.

Vision and the eyes

In an ominous couplet in Shakespeare's *Othello*, Desdemona's father warns Othello:

> Look to her, Moor, if thou hast eyes to see;
> She has deceived her father and may thee.
> (Act 1, scene iii)

It would not have occurred to him to question whether there is any way (literally) to see without eyes (figuratively, Othello cannot see well at all, which causes his downfall). But philosophers must sometimes ask whether what seems patently obvious is in fact true. Let us, then, go a step further than treating light as inessential to seeing.

Suppose Emma has lost her eyes in an accident, but a camera is later connected to her brain in the way her eyes were. When she points it in a given direction in good light, she has just the visual sensations, say of color and shape, that she would have had by looking with her eyes. Might this not be seeing? Indeed, do we not think of the camera as *functioning* like the eye? If, under the right causal conditions, she gets the right sorts of sensations through her eyes or a *functional equivalent* of them, she is seeing.

But are even "eyes" (or organs functioning like eyes) necessary for seeing? What if someone who lacks "eyes" could get visual sensations "matching" the objects in the room by strange radiations they emit? Suppose, for instance, that moving the coin away from the person results in the person's visual impression's representing a decrease in its size, and that the impressions of it are eliminated entirely by enclosing the coin in cardboard. This confirms the presence of an appropriate causal connection between the coin and the discriminative visual experience of it. If no part of the body (other than the brain) is required for the visual impression of the coin, there is no organ plausibly considered a functional equivalent of eyes, but might we not have seeing?

If what is crucial for seeing an object is its producing visual sensations suitably corresponding to it and appropriately responsive to changes in it, presumably the case is one of seeing. If seeing requires the use of an eye or equivalent organ, then it is not—unless the brain itself is a visual organ. It is clear enough that the person would have knowledge of what we might call visual properties, above all colors and shapes. One might call that visual knowledge. But visual knowledge of this kind could be held not to be grounded in seeing, nor acquired through use of any sense organs. For these reasons, we might doubt whether it must be a kind of *perceptual* knowledge. But a case can surely be made for the visual sensation conception of seeing, as against the organ-of-sight conception.

This case, however, may be challenged: can there be "blind sight," seeing in the absence of visual sensations? Something like this is reported in the psychological literature. Imagine an ideal case in which a person with excellent blind sight can navigate among obstacles as if the person saw them, while honestly reporting an absence of visual sensations. Could this be seeing?

We automatically tend to understand such behavior in terms of seeing,

and there is thus an inclination to say that this is seeing. The inclination is even stronger if light's reaching the eyes is necessary for the person to avoid the obstacles. But if the subject has no visual sensations—as opposed to lacking ordinary awareness of such sensations—it is not clear that we must say this, and I doubt that it would be true. The most we must say is that the person seems to *know* where the obstacles are. Knowing through some causal process by which objects produce true beliefs about them is not necessarily perception, and certainly need not be seeing.[11]

It may seem that blind sight is genuine seeing because it produces knowledge of "visual propositions"—propositions ascribing visual properties. But knowledge of these propositions is possible without vision, for instance by something like sonar. Moreover, even dependence on light does not establish that the process in question is visual: the light might somehow stimulate nonvisual mechanisms that convey information about the objects emitting it. Similar questions arise for the importance of sensations to perception in the other sensory modes, for instance of auditory sensations in hearing. There, too, we find hard questions for which competing answers are plausible.

It is difficult, then, to provide an overall philosophical account of just what seeing, or perception in general, is. All the theories we have discussed can help in answering the questions just posed, but none does so in such a simple and decisive way as to leave all its competitors without some plausibility. Still, in exploring those theories we have seen many important points about perception. It is a kind of causal relation. Even its least complex and apparently most basic mode, simple perceiving, requires, in addition to the perceiver, both an object of perception and a sensory experience that in some way corresponds to that object and records, if only imagistically, an indefinite and possibly quite extensive amount of information about the object. Partly on the basis of this information, perception tends to produce beliefs about the perceived object. It implies that the perceiver at least normally has justification for certain beliefs about the object, and it normally produces both justified beliefs about that object and knowledge of it.

Perception may be illusory, as when something appears to have a property it does not have, such as ellipticality when it is really circular. Perception—or sensory experience that seems to the subject just like it—may also be hallucinatory, as with Macbeth's dagger. When it is, the question arises whether there must be interior objects, sense-data, with which perceivers are directly acquainted. But both illusions and hallucinations can apparently be accounted for without positing sense-data, and thus without adding a further element to the four that seem central in perception—the perceiver, the object perceived, the sensory experience, and the causal relation between the object and perceiver in virtue of which that experience occurs. Illusion and hallucination can also be accounted for without denying that perceptual experience—the evidence of the senses—normally yields justified belief and knowledge about the world outside the perceiver. Many questions remain, but so far we have

seen no reason to doubt that perception is a rich and basic source of both knowledge and justification.

Notes

1 The theory of appearing has not been widely defended, but a detailed sympathetic treatment is given in William P. Alston's 'Back to the Theory of Appearing', *Philosophical Perspectives* 13 (1999), 181–203.

2 That the hallucinatory experience is (in its sensory and other phenomenal properties) qualitatively *exactly* similar to the corresponding veridical one, as a proponent of the argument from hallucination intends, is denied by some philosophers—*disjunctivists*—who hold that there are two qualitatively different elements here (hence a disjunction), not a single experiential element common to both cases. Disjunctivists need not deny the plausibility of the argument, however, but evaluating their case against its first premise would require much discussion. For an informative non-technical treatment of disjunctivism see William Fish, *Philosophy of Perception* (New York: Routledge, 2010), esp. Chapter 6.

3 For a recent study and defense of a sense-datum theory see Howard Robinson, *Perception* (London: Routledge, 1994). Cf. Laurence BonJour's *Epistemology: Classical Problems and Contemporary Responses,* 2nd edn (Lanham, MD: Rowman and Littlefield, 2009).

4 The view that ordinary perceptual belief is non-inferential is controversial and—for various senses of 'inference'—has been widely discussed by both philosophers and psychologists. Not *all* sense-datum views, moreover, take perceptual belief to be non-inferential. For a discussion of perception that brings to bear both psychological and philosophical literature see John Heil, *Perception and Cognition* (Berkeley and Los Angeles: University of California Press, 1983), esp. Chapter 2. Cf. Armstrong, *Belief, Truth and Knowledge.*

5 Granted, the mouth of the cup does not appear to us *to be* elliptical if we realize its shape cannot be judged from how it visually appears at an angle, but that is a different point. It concerns what shape we *take* it to have, not what shape visually appears in our consciousness antecedently to our taking it to be of any particular kind.

6 For a detailed and influential discussion of the adverbial theory, with criticism of the sense-datum view, see R.M. Chisholm, *Perceiving* (Ithaca, NY: Cornell University Press, 1957).

7 This is a very important point. One major materialist theory of the mind–body relation—*the identity theory*—says that mental phenomena are identical with brain states or processes. But this theory fails if sense-data exist as mental entities and have properties, such as being green and rectangular, that no brain process has. Identity theorists thus generally oppose the sense-datum theory. See, for example, J.J.C. Smart's much-

discussed 'Sensations and Brain Processes', *Philosophical Review* 68 (1959), 141–56.

8 These and other problems are brought against the sense-datum theory by Winston H.F. Barnes in 'The Myth of Sense-Data', *Proceedings of the Aristotelian Society* 45 (1944–5), 89–117. Cf. R.M. Chisholm's discussion of the problem of the speckled hen: is there, for instance, any answer to the question of how many spots there are in an image of such a hen? And how can we distinguish miscounting the number there are from the number's changing as we count or recount? See his *Theory of Knowledge*, 3rd edn (Englewood Cliffs, NJ: Prentice-Hall, 1989), p. 25.

9 For a detailed twentieth-century defense of phenomenalism, see Book II of C.I. Lewis's *An Analysis of Knowledge and Valuation* (La Salle, IL: Open Court, 1946). Cf. R.M. Chisholm's widely known criticism of this defense in 'The Problem of Empiricism', *Journal of Philosophy* 45 (1948), 512–17.

10 Berkeley might hold that if *God* has booklike sense-data, it does follow that there really is a book. A case can be made for this, but one might also argue that as an all-powerful being God could bring it about that there is a distinction between his creating a physical object and having the corresponding sense-data.

11 A subject who really *does* have visual impressions could also misreport. The possibility of such misreporting about one's own consciousness is discussed in Chapter 3.

3 Memory

The preservation and reconstruction of the past

- **Memory and the past**
- **The causal basis of memory beliefs**
- **Theories of memory**
  Three modes of memory
  The direct realist view
  The representative theory of memory
  Memory images
  Remembering
  The phenomenalist conception of memory
  The adverbial conception of memory
- **Remembering, recalling, and imaging**
- **Remembering, imaging, and recognition**
- **The epistemological centrality of memory**
  Remembering, knowing, and being justified
  Memorial justification and memorial knowledge
  Memory as a retentional and generative source

3 Memory

The preservation and reconstruction of the past

I believe that I have pruned the flowering crab apple tree which stands in the center of the lawn. This belief is apparently grounded in my memory. When I look at the tree and notice its shape, it often occurs to me that I have pruned it. When this does occur to me, I have a sense of already believing it. The proposition that I have pruned the tree does not seem to be a discovery or a result of inference or a bit of wishful thinking, but rather something I have had in mind before and now believe with some conviction.

On the basis of all these facts about my belief that I have pruned this tree—especially my sense of having already believed this—the belief is justified. Indeed, I cannot help also thinking I know that I have pruned the tree. In particular, my belief that I have seems to be grounded in memory, in the way that what I genuinely *remember* is grounded there. Consider remembering one's having just read the preceding part of this page. If one has just done so, there is likely to be a clear sense of having done it. We do not, for instance, simply have a dreamlike recollection, nor are we concluding what we seem to remember from something else, as you might conclude, from the distinctive shape of a tree, that it must have been you who pruned it.

What, in general terms, is memory? Is it anything beyond a storehouse of some of what we have experienced and learned? And what is it to remember something? Whatever remembering is, its objects include people, material things, facts, events, and, among the events, our own experiences. We might also say that remembering is the chief "function" of memory. Is remembering, then, exercising, or being able to exercise, the capacity of memory? And is there—as with perceptual knowledge by contrast with mere perceptual belief—a special kind of *success* that goes with remembering something as opposed to simply believing it from memory?

In pursuing these questions, it is useful to compare memory with perception. Both are crucial for our knowledge of what is external to the mind: the latter gives us a view of what is outside of us in the present, the former of what is outside of the present altogether. Moreover, memory builds on perception; it preserves much important information we acquire through the senses. It also preserves information about our mental lives. But how does memory achieve this preservation? Must it, for instance, operate by storing

images, or can it preserve bare facts? Before we can see how memory is connected with knowledge and justification, we must first understand what it is and something about how it works.

Memory and the past

We can learn some basic points about memory and remembering by clearing away some tempting mistakes. To begin with, we cannot say simply that memory is a capacity for knowledge or belief about the past. It is true that memory entails that capacity; but one could have and even exercise the capacity *without* exhibiting memory of the past. Consider the events of World War II. I can know a good deal about them through reading, but at the time I am learning about them through reading I have no memory of them. I witnessed none of them, and I do not remember them. To be sure, I may remember a *description* of them and thereby say—perhaps recalling a history class—that I remember (for instance) the invasion of Normandy. This could be called remembering the events *indirectly*. But it is not remembering in the direct and primary sense that concerns us.

Far from all knowledge of the past being a kind of remembering, then, we commonly know propositions about the past on a basis other than remembering it. Consider again the knowledge of the past obtained while reading; this knowledge is not a case of *remembering* the past but a kind of knowledge of the past acquired through testimony about it. Similarly, I can gain knowledge about the past from your present description of what you did yesterday. This knowledge may not be *retained*, hence need not become memorial. It may never get into the storehouse: I could lose it after I have acquired it, just as we forget a phone number needed only for a moment. In these instances, I have knowledge of the past, but only for too brief a time to qualify as *remembering* the propositions I momentarily knew.

The same example shows a second major point. Like knowledge of the past, beliefs about the past, such as those I acquire about your activities, do not necessarily represent memory. For they need not be *retained* and so are not *memory beliefs*, that is, beliefs grounded in the "faculty" of memory. They are grounded in testimony and are forgotten before being memorially stored.

Moreover, even when one does memorially retain beliefs about the past, they need not amount to *remembering* something. Retained beliefs about the past can be sheer fabrications unconnected with memory capacities. Imagine, for instance, that although I have not seen you for a year, for some reason I groundlessly form the belief that precisely a month ago you wore the belt I see you wearing now. This belief is not memorial, and even when retained, it would not constitute remembering. It comes not from memory of a past event but from undisciplined imagination. Retention of a conviction grounded in fantasy does not upgrade it into remembering.

One might think that beliefs about the past, when they *are* memorial, and not merely retained, constitute remembering. But this need not be so,

because they may be false, whereas everything we genuinely remember to be the case is true. Remembering is, then, *factive*. If, for instance, I remember that Thomas Reid discussed John Locke's ideas about memory, then he in fact did.

To see that even a vivid memory belief that something is so need not represent genuinely remembering that it is, suppose my memory plays a trick on me and I misremember an actual event. I mistakenly think, and vividly imagine, that I planted a green spruce when it was really a blue one I planted. I would now have a memory belief (one tracing back in a normal way to an event it is about) that is mistaken, even though its close relatives in my memory are true. Still, I cannot remember that I planted a green spruce if in fact I did not. Even when a falsehood resides in the neighborhood of truth, it is not elevated into an object of remembering simply because it is retained in a memory belief.

Suppose, however, that a retained belief about the past is true. Is it then an instance of remembering? Not necessarily. Even true beliefs about the past may be utterly baseless and true only by lucky accident. Suppose that my memorially retained belief that you wore that belt just happens to be true because by chance you did select the same belt for both occasions. This belief still does not represent *remembering*. I have merely retained my luckily true impression that you wore the same belt. A retained belief of this sort is stored in memory, but only *properly grounded* true beliefs stored there constitute remembering that something is so.[1]

The causal basis of memory beliefs

One might think that just as perceptual beliefs are caused by an object perceived, memory beliefs are caused by a past event remembered. Some memory beliefs are caused in this way, and we will soon see that causal connections to the past are essential for genuine remembering. But even if it should be true that all memory beliefs are produced at least partly by events in the past, past events are not the only *objects* of memory or the only things it "stores." We remember, and thereby retain and believe, general truths, such as mathematical theorems. Mathematical propositions are certainly not past events (propositions are not events of any kind). *Learning* them is a past event for many people, but that is a quite different point. Nor are the propositions past objects of some other sort, or even about the past; but many truths of mathematics are clearly among the objects of remembering—the things we remember.

Moreover, even if every memory belief is at least partly caused by a past event, not every belief caused by a past event is memorial. This point applies even if the belief is true. Suppose that my unknowingly taking a poisonous drug causes me to feel strangely ill an hour later, and my feeling ill then causes me to believe I have been poisoned. Then, indirectly, the past event of taking

poison causes me to believe that I have been poisoned. But this belief is not memorial: it is in no way grounded in my capacity for remembering, and I have no memories connected with the belief, such as a memory of someone's putting a white powder into my soup. I arrive at the belief by inference to what I think best explains my illness. Thus, the belief's being caused by the past event of my taking the drug need not make it a memory belief, even if the belief is true. My memory has played no role in supporting the *content* of the belief. The belief lacks a ground appropriate for suitably connecting it with the past event it represents.

An analogy with perception will help. Consider a belief caused by a flash that I do not see, but merely feel as a momentary heat. This belief need not be a visual belief, even if it is a true belief with visual content, say that a camera flashed near my hand. A belief caused by something visible is not thereby a visual belief, just as a belief caused by a past event—something remember-able—is not thereby a memory belief. Since a belief caused by a past event need not be memorial, it is not true that a memory belief simply is a belief at least partly caused by a past event.[2]

The analogy between memory and perception is limited, but it does get us on the right track. For surely a belief about the past is memorial *only* if it has some causal connection to a past event, just as a belief is perceptual (say, visual) only if there is some causal connection between it and the perceived object. Even a belief that arises from testimony and not from first-hand observation and is then stored in memory is *traceable* to the past event of one's acquiring the belief. A thing cannot normally be stored in memory unless it has entered that storehouse. Since memory beliefs can concern any subject, including future events or mathematical truths, we can see that such beliefs need not be *about* any particular event even if their existence does *trace* to one.

Could we, however, have *innate* beliefs? If so, they could be about the past but not memorially connected with a past event, perhaps because the belief is possessed at the time one came into being and does not trace to any remembered experience. It would not *enter* the storehouse of memory: it is part of one's initial equipment. To be sure, perhaps an innate belief could be memorial in roughly the ordinary way if in some *previous* incarnation it traces to an appropriate event, something of the kind that produces a memory.[3] Otherwise, it is best considered merely a *retained belief*, say retained from birth as part of one's native endowment, rather than a belief entering one's memory through, say, observation or testimony.

Just as it is hard to specify how, in order to be perceptual, a belief must be causally connected to the perceptible object it is about, it is hard to specify how, in order to be memorial, a belief must be causally connected to the past. This will become clearer as we explore memory, but fortunately many points can be made about memory without a detailed account of the kind of causal connection in question.

Theories of memory

If we view theories of memory on the model of three major kinds of theories of perception discussed in Chapter 1, there is much we can discover both about the kinds of causal relations required for remembering and about how memory grounds justification and knowledge. Broadly speaking, the three kinds are direct realism (including the adverbial theory and the theory of appearing as well as naive realism), representative realism, and phenomenalism. Each has an analogue in the theory of memory.

Three modes of memory

In constructing theories of memory, there are at least three different but closely related notions we must track: *memory*, *remembering*, and *recalling*. We have memories of many things. We remember, and we sometimes recall—roughly, call back to mind—much that we have experienced. Both points apply to us in virtue of the power of our memory, conceived as a capacity (a mental "faculty"). There are things we remember, such as isolated facts, that we may never have occasion to recall. But they remain in the storehouse of memory ready to be retrieved if needed. When retrieved, we may be said to have, at the time of retrieval, *a memory* of them.

Our memory, conceived as a "faculty," is a general capacity: the better it is, the more memories we can have, the better we remember, and the more we can recall. Among the things we remember are *skills* and related behavioral capacities, both mental and physical. Memory of skills is *remembering how*. Much of what emerges here concerning remembering will apply to remembering how (though the notion does not seem reducible to any kind of knowing *that*), but I will not take time to discuss it specifically.

There are, then, three memorial notions to be accounted for by a theory of memory: first, remembering of events, things, and propositions; second, recalling those items; and third, memory as the capacity in virtue of which remembering and recalling occur. There is a further task: accounting for errors. Like perception, memory, as the capacity *for* remembering and recalling—and I include *recollection* as a kind of recalling—can produce impressions that are illusory or, in a way, hallucinatory. Not every memory belief is true; not every recollection is faithful to what it recalls.

In developing the memorial counterparts of the three main kinds of theories of perception, I will concentrate chiefly on remembering, particularly on the simple remembering of events—*event memory*, for instance of my pruning the tree—as opposed to remembering *that* I pruned it, *propositional remembering*, or remembering the pruning *to be* hard, *objectual remembering*. I assume that, like simple perception of something, simple remembering of an event, such as a bird's flying by, does not entail having a belief about it, as opposed to being disposed to form beliefs about it if the occasion elicits them. But let us concentrate on cases in which one does have such a belief. These cases are crucial for understanding memory knowledge.

The direct realist view

There is a memorial counterpart of naive realism in the theory of perception. It is the view that when we remember an event, we just plain remember it and it is as it seems to us to be. This might be taken to mean that the event is directly presented to us by our memory, as if it were *present* in memory, just as a flash might be present to vision. The difference is that unlike a flash that fills one's visual field, the remembered event is not taken to be occurring. Like all the plausible accounts of memory, this one is best construed as causal: as assuming that some causal chain links us to the remembered event. If I remember seeing Bill a year ago, then it must be in part *because* I did see him that I believe (or am disposed to believe) that I did, and not, say, because I dreamt that I did.

As a direct realist view, this position also maintains that a memory belief is not produced by any intermediary with which we are acquainted, such as an image. To say that would imply a counterpart of the sense-datum theory. We would have an *in*direct realism: just as we perceive the outside world through sense-data that present it to us, we remember the past by virtue of memory's presenting it to us.

Naive realism about memory is inadequate. To begin with, not just any causal connection to the past will do, as we saw with the poisoning case. The causal chain linking a memory belief to a remembered event must be in a sense *unbroken*. In part, the idea is that a belief retained in memory cannot be lost from it during the period of retention. To see the idea, consider a broken chain. Imagine that you saw me prune the apple tree and you remember my doing so. The pruning is then the main causal ground of your memory belief, as it is of mine, and we both remember my pruning it. But suppose I completely forget the event and thus no longer believe I pruned the tree, *then*, solely on the basis of your testimony, later come to believe (once again) that I pruned it. There is still a causal chain from my present belief back to the pruning; for the pruning produced your belief, which in a way produced your testimony, which in turn produced my present belief. But the memorial chain in me was broken by my forgetting.

Given this kind of broken chain to my pruning, I do not retain my original belief and do not remember my pruning; I simply know, from your testimony, that I did it. I now believe the "same thing" but do not have the same (the original) belief.[4] To be sure, *after* your testimony, when I have retained the knowledge you gave me, we might say that I now remember, again, *that* I pruned the tree but no longer remember *pruning it*. Propositional memory *about* an event, even an action of one's own, does not entail event memory *of* it.

The case also shows that knowledge of a past event, even if it is one's own action, does not entail remembering it. I know *that* I pruned the tree, but I do not *remember* pruning it. My propositional knowledge of the event no more represents remembering *it* than my knowledge based solely on your testimony that there is a radiant sunset visible from the front porch represents my

seeing it, when I am inside reading. If, however, I know that I pruned it, then that event is a real element in the past.

A realist view of remembering seems correct, then, if it is coupled with the requirement of an unbroken causal chain. But as stated so far, the view is deficient in some of the ways that naive realism about perception is. For one thing, memory is subject to illusion. I might remember an event, such as meeting you, but not quite as it was, just as you might see white paper in yellow light, and thus not see it as white but as yellow. Here I do not simply remember; I remember incorrectly, for example in remembering the meeting as taking place in New York when it was in fact in Chicago. (I correctly remember meeting you; I do not remember the location of the meeting.) Second, there is the memorial counterpart of hallucination: I may have a vivid image of mailing a letter, and might believe I remember doing so, yet be mistaken. We must, then, account for memorial illusion and similar problems.

The representative theory of memory

The territory may begin to look familiar, particularly if we recall the sense-datum theory of perception, which posits inner sensory objects that, as intermediaries between external things and the mind, represent the former to the latter. For instance, suppose that there are memory images, and that they are genuine objects which figure in remembering rather as sense-data are thought to figure in perceiving. These images might even *be* sense-data if they are vivid enough, but normally they are more like the images of imagination. It might be like this: seeing the apple tree as I prune it produces sensory images in me (whether these are sense-data or not); my memory images of the tree might be conceived as a kind of *residue of perception*.[5]

Perhaps, then, we may be said to remember an event when we have at least one true belief about it suitably grounded in a memory image of it, that is, an image of it which derives, by a suitable unbroken chain, from our experience of the event and represents it correctly in at least some way. The better the memorial representation of the event, the better our memory of it. Call this view *the representative theory of event memory*. It takes event memory to be a representational faculty that works through images that "picture" what they represent.

Memory images

Like the sense-datum theory of perception, the representative theory of memory is an indirect realism. It construes our remembering as *mediated* by memory images (though not as based on inference from facts about such images); it is *through* images that we are acquainted with the past. The view is also like the sense-datum theory in readily accounting for memorial illusion

and similar problems. To remember incorrectly, as opposed to simply having a false belief about the past with no basis in memory, is to be acquainted with a memory image that, despite its being sufficiently faithful to the remembered event to ground one's remembering it, has *some* aspect which produces a false belief about the event, say that it was in New York rather than in Chicago.

The counterpart of hallucination is also treated as one would expect by analogy with the sense-datum theory. Memorial hallucination occurs when one has an image that is intrinsically like a memorial one, but not linked to a past event by a suitable causal chain, just as, in perceptual hallucinations, the sense-data are not produced by the object (or are produced by it in an abnormal way).

Unfortunately, the representative theory of memory has many of the difficulties of the sense-datum theory and some of its own. Consider the similar difficulties first, particularly in relation to remembering.

Remembering

Remembering an event surely does not require acquaintance with an image of it. You may be able to reel off, from memory, some details of a conversation you heard a week ago, even if you have no images, even auditory ones, of the conversation or what it concerned. Moreover, misremembering an event does not require acquaintance with something, such as an image, which actually *has* the property one mistakenly remembers the event as having had, as a sense-datum representing the mouth of a cup from a certain angle is supposed to have the property of ellipticality. I can misremember my meeting you by remembering our meeting as being in New York, when it was actually in Chicago, even if the mistaken element in my memory is not accompanied by images of anything in New York (and even if the correct aspect of my memory is not accompanied by an image that is of our actual meeting in Chicago). I may simply remember the occasion with its animated conversation, yet have the false impression that it was in New York.

Memorial thinking—an episode of thinking about one or more remembered objects or events—may also be possible without objects to serve as images of the past. In retrospective imagination, might I not vividly experience our conversation even if I am acquainted with no object that represents it for me in the way that, in hallucinations, sense-data are supposed to represent physical objects?

Granted, if I have no images, then I cannot *recall*—in the sense of bringing back into my visual consciousness—the color of your sweater. But I might still remember what you said and the hoarseness of your voice owing to the flu, and I might remember what color your sweater was even if I cannot bring the color itself to mind (perhaps you said that its pale blue matched your jacket, and by that remark I remember what the color was without imaging it). I can apparently imagine past events—whether accurately or

inaccurately—without having direct acquaintance with memorial pictures of them, just as I can apparently hallucinate an object without having direct acquaintance with a sense-datum representation of it.

A further difficulty for the representative theory arises when we consider a disanalogy between remembering and perceiving. I can remember our meeting and describe it to someone from memory even if I have *no* images or image-like experiences at all, whereas I apparently cannot see a tree if I have no visual sensations, such as the impressions of foliage that make up an image of a tree. Remembering, even of events that one has perceived, is neither a sensory event nor necessarily an *imaginational* one (even if it often is, especially in some people, such as those who are highly "visual"). So there need not be, in every case of remembering, even the *makings* of a representative theory to which images are crucial.

The phenomenalist conception of memory

The kinds of difficulties we have seen in relation to the representative theory of memory suggest that the memorial counterpart of phenomenalism may also suffer irreparable difficulties. Above all, a phenomenalist account of memory relies on images or imaging at least as heavily as does the representative theory, but neither images nor imaging seems either necessary or sufficient for remembering events. Let us explore this.

On what may be the most plausible phenomenalist account of memory, remembering an event is understood in terms of the imaginational content of present experience. To remember an event is (roughly) to have a suitable collection of images representing it, on the basis of which, in a certain way, one believes (or is disposed to believe) something about that event.

This will not do. Remembering an event simply does not require a collection of images analogous to the sense-data from which phenomenalists try to construct physical objects (or even a collection of imaging experiences such as an adverbial phenomenalist might posit). Images of the kind posited to account for remembering are not only not necessary for remembering, as our examples show; they also are not sufficient for it either. Just as no collection of sense-data is such that its existence implies perception of an external object, no collection of images (even apparently memorial images) is such that, in having a belief about the past grounded on those images, one must be remembering something. No matter how vivid my images of talking with you beneath the skyscrapers of Wall Street, I might not remember our talking there, and my belief that we did talk there (or anywhere) might still be mistaken.

The adverbial conception of memory

If these difficulties are as serious as they seem, then if, in search of a better account of memory, we are to change course and construct a plausible

alternative theory of remembering, we must take account of them. First, such a theory will not claim for remembering all the kinds of directness it posits for perception. Plainly, memory is not *temporally direct*, as past events are not temporally present, whereas we can see a thing's properties at the same time that it has them.[6]

By contrast, any plausible account of remembering, such as a well constructed adverbial theory of it, will take remembering to be (as perceiving apparently is) *epistemically direct*. Memory beliefs, as we have seen, are not inferential. It is not on the basis of any *premise* that I believe (or know) that I have pruned that crab apple tree. My belief is grounded in memory as a preserver of beliefs and other elements, not in other beliefs which give me premises to support the memory belief.[7]

Moreover, an adequate theory must not say that (actively) remembering an event, such as pruning a tree, is constituted by memorially imaging in a way suitably caused by that past event, as perceiving an object is sensory stimulation suitably caused by the thing perceived. For no such imaging need occur (though it commonly does, especially in the case of the active remembering that constitutes recalling). We can describe a past event to others, and in doing so actively remember it, even if we are imaging nothing but their faces.

Positively, the *adverbial view of memory*, applied to remembering events, should be expressed as something like this. First, *actively* (*occurrently*) *remembering* an event is *manifesting—realizing,* in a sense—a memorial capacity concerning it, in which this manifestation is linked to the event by an unbroken causal chain. Just as, in observing a cat, one is manifesting a perceptual capacity, in describing a play from memory one is realizing a memorial one. The most typical manifestations—the things that constitute experiencing in a memorial way—are probably (1) imaging processes concerning the event; (2) formations of memory beliefs about it, often through considering one's memory images; and (3) considering the propositions so believed, with a sense of already believing them. But there may be other realizations, for instance recognizing a picture of the event. Second, *passively* (*dispositionally*) *remembering* an event is having this capacity in an unmanifested way, as when, though I can recall the pruning if I want to, my mind is wholly on other things. For me to remember the pruning actively, something must call it to mind.

To see the difference between the dispositional and the occurrent in another context, consider elasticity in a rubber band. It is a dispositional property, whereas stretching is an occurrent property that manifests the disposition (this distinction is further discussed in Chapter 4). Recalling an event can be related to dispositionally remembering it much as a thing's stretching is to its elasticity. Just as stretching manifests the disposition of elasticity, recalling is a case of actively remembering that manifests the dispositional memory which retains the thing recalled.

Propositional remembering—remembering *that*—can be construed similarly. On the adverbial view imagined, to remember that an event occurred

is a memorial *way* of truly believing that it did, roughly, to have one or more true beliefs about it which are suitably linked by an unbroken chain to past experience and represent the event in a certain way (if only as occurring). These beliefs constitute knowledge that is preserved in memory. On this view, then, remembering that something is so constitutes knowledge *from memory*, rather as seeing that the cat is sleeping constitutes knowledge through perception.

Most of what we propositionally remember is dispositional, roughly, recorded in dispositional memory beliefs. When these beliefs are called up in active propositional remembering, as when I describe how I pruned the tree, one is experiencing in a memorial *way*. This does not require being acquainted with imagistic memorial objects. One may, but need not, image memorially, as when one actually calls up the remembered experience and focuses on its features in one's imagination.

Moreover, whether one images a remembered event or not, the event need not be entirely *as* one remembers it. Here event memory differs from propositional memory; the former, like seeing *as*, can misrepresent the thing in question, whereas remembering *that* something is so, like seeing that it is, entails its being (exactly) so. One can remember a meeting as being in the wrong city, thus remember it in the wrong way geographically, just as one can see a circular cup as elliptical and so see it in the wrong way visually.

In neither case of illusion, to be sure, does one have to be fooled: with memory as with perception, illusion does not always produce false belief. Typically, if I remember something as having a certain quality, say a conversation as being rushed, I believe it was like that; but I can remember it *as* such, yet know from independent evidence (such as testimony) that it was not rushed. If, however, we really remember some object or event, then we are right about some aspect of it, or are at least in a position to form some true beliefs about it on considering the matter. This is parallel to the point that if one really sees something, one is at least in a position to see it to *be* something or other.

Remembering, recalling, and imaging

So far, the adverbial view seems superior to its competitors in relation to the crucial notions to be accounted for, such as remembering and recalling. Will this direct realist view stand scrutiny? In answering this, it is important to see that the view can account for imaging; it simply does not take imaging to be an acquaintance with inner objects. Still, there may be a nagging doubt about whether it does not incline us to posit too little imaging. When I am remembering an event, especially a perceived one like a ship's docking as opposed to an imperceptible one like thinking about knowledge, I typically *do* image some aspect of it. I refer, of course, to active remembering, as opposed to my stored remembering of events that are now far from my mind but which I *could* actively recall if the subject came up. The first kind of

remembering is occurrent, as it is in part a matter of something's occurring in me. The second kind is dispositional, as it is a matter of my being disposed (roughly, tending) to remember a thing actively (occurrently) *provided that* something activates my memory. Thus, although yesterday's concert may be far from my mind while I write a letter, if someone asks how I liked the Chopin, then my dispositional memory may be activated; and, as I recall it, thereby occurrently remembering it, I may say I thought it inspiring.

It is occurrent remembering that is analogous to perception and is of chief concern now; and it is occurrent remembering that is closely associated with imaging. Does occurrent remembering require some sort of imagery after all, even if not images as sense-datum objects?

Here is a natural way to answer. Consider one of your memories of an event, for instance meeting someone for the first time. Do this in such a way that you take yourself to be actively remembering that event. Second, ask yourself whether you are now imaging. When I do this, I image. Here, remembering involves imaging. But notice what has happened: I have *called up* a memory and inspected the results of my effort. Perhaps I am imaging because of the *way* I evoked the remembering, or because I scrutinized the process of my calling up the meeting. Self-conscious evocation of the past and scrutiny of the results may yield findings unrepresentative of remembering in general.

This procedure of evoking memories of the past, then—selecting them by recalling past phenomena—is defective as a way of determining whether remembering requires imaging. But the procedure does show something. For suppose that what I have done is to *recall* a past event. Perhaps recalling, which is calling back to mind, often by a lengthy search of one's memory, *does* require imaging provided it is a recalling *of* an imageable event, such as pruning a tree, as opposed to, say, recalling a theorem. There is some reason to think this. If no imaging of our seaside luncheon enters my consciousness, how can I have recalled it? Sometimes, moreover, we say that we cannot recall someone, meaning not that we do not *know* who the person is, but that we cannot image the person. There, recalling seems to imply imaging. To be sure, recalling the content of our conversation is possible simply by becoming conscious, in a certain recollective way, of its informational content. But the content of a conversation, as distinct from our activity of conversing, is not imageable in the first place.

Even if the most common kinds of recalling should imply imaging, however, remembering does not. Why, then, does that idea persist? For one thing, when we collect specimen memories in order to examine remembering, we often do it *by* recalling things. If so, it should be no surprise that the specimen memories involve recalling something. If, in trying to determine the shades of beech leaves, I collect specimens only from the nearby copper beeches, it is no surprise that I may erroneously think beeches in general have copper-colored leaves.

Remembering, imaging, and recognition

A related point is that what we cannot recall we often *believe* we cannot remember. On the adverbial view, this is natural; for an inability to remember is a lack of a memorial capacity, and, understandably, we may think we lack that capacity when, under normal conditions, we cannot exercise it in an expectable way—such as recalling an event we have been taking ourselves to remember. But imaging is only *one* exercise of memorial capacity, important though it is; and just as we can be capable of climbing a mountain, but not necessarily by every route to the top, we may have the capacity constituted by remembering something, but be unable to exercise the capacity in every way it can be exercised. Hence, inability to image does not imply that one does not remember the thing in question. We can see, then, both why there is a tendency to think that remembering requires imaging and why we should not accept this view.

Imaging may still seem more important for remembering than so far granted. But take another case. Suppose I can neither recall nor image Jane. I might still remember her; for on seeing her, I might *recognize* her, and might even recall, our last meeting. This would suggest that my memory simply needed to be "jogged." In adverbial terms, before I see her again I dispositionally know her in a certain memorial way—I remember her—even though I cannot imagistically experience her in a memorial way—namely, recall her.

I choose the example of remembering a person because it is easy to show that we do remember someone by creating the right occasion. Recalling her is an indication of my remembering her, but it may not be possible despite my remembering her; recognizing her when I meet her is a proof of the pudding. We cannot draw this contrast with past *events*, since, unlike people, they cannot be literally brought back. But even here, there is indirect recognition, as when one recognizes a ship's docking in Helsinki harbor upon seeing a picture of the event. It is doubtful, then, that the relation between recalling and remembering is different with events.

It is important to see that the way I am now considering the relation between recalling and remembering is direct and non-introspective. I am exploring what is possible and what it would show. It is possible, however unlikely, that I might have no retained image of pruning the apple tree, yet be able to give an account of the pruning that is both remarkably accurate *and* grounded by a suitable causal chain in the original experience of the pruning. If I do this without my having received any information about the event from anyone else, it is an excellent reason to think I remember the event. It is akin to recognition of a person one could not recall, say by picking the person out of a crowd.

To be sure, our *beliefs* about what events we remember may depend on what we can recall, which, in turn, may be largely dependent on what we can image. But what events we *do* remember is a matter of how our memorial capacities are grounded in the past and not of what kind of evidence we can get, imagistically or otherwise, concerning that grounding.

In exercising my capacity to remember events, then, I need not rely on my images or even on my ability to image, though in fact retention of images doubtless aids remembering. The representative theory of memory therefore seems mistaken, and some memorial analogue of direct realism regarding perception is apparently preferable. The possibility of a good analogy is already implicit in the point that an event can be perceived even though the time at which one has a sensory experience representing it is later than its occurrence. The suggested adverbial view of remembering is a good position from which to work; but I leave some important questions about memory unexplored, and it would be premature to present that view as clearly correct.

The epistemological centrality of memory

We can now see some points about memory as a source of belief, knowledge, and justification. Let us start with beliefs. Memory is a source of beliefs in the way a storehouse is a source of what has been put there, but it is not a source of beliefs in the generative way perception is. Clearly our memory, as a mental capacity, is a source of beliefs in the sense that it *preserves* them and enables us to *call them up*. It also enables us to *draw on our beliefs to supply premises* in reasoning. We do this when we solve mathematical problems using memorized theorems. We may also be guided by other kinds of presupposed premises without having to call them to mind. Remembered propositions (and patterns) can be like routes we know well: they can guide our journey while we concentrate on the road just ahead.

When our memory beliefs are of propositions we remember to be true, they constitute knowledge. If you remember that we met, you know that we did. Similarly, if you remember me, you know me, at least as I was once—you may not be able to recognize me now. So memory, when it is a source of what is remembered, commonly yields both knowledge *that* and knowledge *of*. Remembering, then, is *knowledge-entailing*. The analogy to perception is significant here too.

Is memory also a source of justification? Surely what justifies the great majority of my justified beliefs about the past is my memory. For instance, my belief that I twice pruned the crab apple tree is justified because of the way that belief resides in my memory. It has, for example, a special kind of familiarity, confidence, and connection with other things I seem to remember. Moreover, it appears that if I remember that I met you, I am justified in believing I met you. It thus seems that when memory yields genuine remembering it also yields justification. Certainly this commonly holds.

Remembering, knowing, and being justified

Perhaps, however, I could remember that I met you, yet *fail* to be justified in my belief because (in fun) you convince me, by good arguments and by enlisting the corroboration of plausible cohorts, that I am probably confusing you

with someone else. Still, suppose my belief remains properly grounded in my actual memory of having met you, perhaps because the memory is so clear that the belief is almost unshakable. Then I may still genuinely remember that I met you. Despite this point, if your arguments are good enough, I may properly reproach myself for still holding the belief that I met you, and my belief may perhaps cease to be justified. Its justification would be defeated by your arguments and by my own credible self-reproaches based on seeing their plausibility.

If this case is possible, it has an important implication. If, as I have suggested, remembering that something is so entails knowing it is so, then the case as described implies that knowing that something is so does *not* imply justifiedly believing it. (In Chapters 11 and 12, I return to the relation between knowledge and justification, but it is important here to see that the domain of memory provides a challenge to understanding that relation.)[8]

Furthermore, if the case is possible and one *can* remember that something is so, yet fail to be justified in one's believing that it is so, then we might question whether memory yields any justified beliefs after all. Fortunately, the example by no means rules this out. Quite apart from cases of genuine remembering, memory often yields justified belief. If I have a vivid and confident belief that I met Jane, and this belief seems to me to arise from a memory of the occasion, I may, simply on that basis, be justified in the belief. Surely this is, after all, just the sort of belief that usually does represent remembering; in any case, I have no reason to question its credentials.

Memory can justify a belief even when that belief does not constitute knowledge or rest on actual remembering of the proposition or event in question. If, for instance, I do not in fact remember meeting Jane, perhaps the only reason why I do not is that it was her identical twin, of whose existence I had no idea, whom I met. That excusable ignorance may prevent my *knowing* that I met Jane, but it does not preclude my justifiedly believing that I did.

Memorial justification and memorial knowledge

These reflections suggest *a memorial justification principle for events: normally, if one has a clear and confident memory belief that one experienced a given thing, then the belief is justified*. We might also call clear, confident memory beliefs prima facie justified.[9] A memory belief is one grounded in memory; this is commonly a kind of belief which represents the event or proposition in question as familiar in a certain way. Commonly, if one considered the matter, the belief would seem to one to arise from one's memory; but the notion of a memory belief cannot be defined by that normal property of such beliefs, and it is not easily defined at all.[10] We can believe—even know—from memory propositions we do not find familiar (as when we have not recalled them or thought of them in years).

A still broader principle may perhaps be true—*the general memorial justification principle: normally, clear, confident memory beliefs with any subject*

matter are prima facie justified. Moreover, if they do not conflict with other beliefs one holds, say that one has never been to the country where one now seems to remember going to a museum, they tend to be justified on balance. With both principles the degree of justification may not be great, particularly if there is no corroboration, such as apparently recalling a sequence of events related to the belief. My belief that I met someone at a restaurant tends to be better justified if I apparently remember related events, such as a friend's recently mentioning our meeting that person there, than if it is isolated from other apparent memory beliefs that confirm it.

Both these and similar principles help to describe how memory is plausibly conceived as a source of justification. This is certainly how it is standardly conceived. Imagine someone saying "I have a clear and confident memory belief that we met at the Café Rouge, but this gives me no justification whatever for thinking so." We can understand someone's holding that there is *better* justification for not believing this—say, because of known memory failure—but that would show only that the justification is defeated, not that there is none whatever to be defeated.

Memory as a retentional and generative source

There is a very important difference between the way in which memory is a source of knowledge and the way in which it is a source of justification. To see this, we must take account of several points. Memory is a *preservative* capacity with respect to both belief and knowledge. First, when you initially come to believe something, you do not (yet) remember it. Second, you cannot (propositionally) remember something unless you *previously* knew or at least believed it, for instance perceptually, and your belief of it is suitably preserved.

Thus, memory *retains* belief and knowledge. Retention is roughly equivalent to preservation but has a lesser implication of unchangingness; a belief held with considerably less confidence than originally, for example, is less properly said to be preserved than to be retained. Memory does not generate belief and knowledge, except in the sense that, by *using* what you have in memory, you can acquire beliefs and knowledge through inference (or perhaps through other processes that themselves yield belief and knowledge). I may, for instance, infer much from propositions I remember, or I may arrive at greater knowledge of a movie I saw by calling up images of various scenes. Here it is thought processes—inferential and recollective—that, *partly* on the basis of retained material, produce belief and knowledge.

To say that memory is not a generative source of knowledge is not to deny that memory is sufficiently connected with knowledge to figure in a plausible epistemic principle—call it the *memorial knowledge principle*: *normally, a true memory belief, supported by a vivid, steady experience of recall that is in turn corroborated by other memory experiences, represents knowledge*. But if this principle is correct, that is *because* such beliefs are of a kind that ordinarily

constitutes knowledge originally, say when one learned through perception the truth that a tree has recently been pruned, and continued therefore to have grounds, preserved by one's memory, for holding this belief.

Memory is not, then, a *basic source* of belief or knowledge, a source that generates them other than through dependence on a contribution by some different source of them. It is, however, a basic source of justification. We can be justified in believing something either on the basis of remembering that it is so, or on the basis of our having a clear and confident memory belief that it is, even if this belief is not true. If, however, we genuinely remember that something is so, it is so, and we know that it is.

Memory can not only generate justification, as when my vivid sense of remembering a line of poetry justifies believing that I do. Memory can also preserve justification, particularly when the justification resides in memorially retained beliefs, for instance beliefs of one's premises for a view one holds. But the original justification of a belief *need* not be retained in order for the belief to be memorially justified. The sense of memory can generate justification by virtue of the way the proposition or event in question occurs to us, and this could occur when we have forgotten our original, remembered justification, such our as witnessing the event that, perhaps because we have read a vivid account of it, we now have a memorial sense of remembering.[11]

This justifying capacity of memory often operates even when we have no associated images. But in accounting for what justifies memory beliefs, images do have a significant if restricted role. We are better justified in a memory belief supported by imagery, especially vivid imagery, than in memory beliefs not thus supported (other things being equal). Perhaps the reason is that we have at least some justification for believing that there is less likelihood of error if both imagery and beliefs point in the same direction, say to my having met you two years ago.[12]

For all the analogy between memory and perception, then, there are important differences. Although both are essential to our justification for believing a huge proportion of what we believe, perception is more fundamental in a way that is crucial to the development of our outlook on the world. It supplies memory with much of its raw material, whereas memory, though it guides us in seeking what to observe and, in that way, often determines *what* we perceive, does not supply raw materials to perception: it manufactures no perceptibles. It does, to be sure, supply raw materials for introspection and thought: we would have vastly less to "look in on" or think about if we did not remember sights and sounds, conversations and embraces, ideas and plans.

Both memory and perception, however, are to be causally conceived, and both are, in different ways, sources of belief, justification, and knowledge, propositional as well as objectual. But perception is a basic source of all three: it can produce them without dependence on contributions from another belief-producing capacity, such as reasoning. Memory, being a capacity for

the preservation, and not the creation, of belief and knowledge, is not a basic source of them.

Still, without memory, perceptual knowledge could not be amassed and used to help us build theories of the world or of human experience, or even to make local maps to guide daily living. We would not even have a sense of who we are, since each moment in our lives would be dead to us by the next. Beyond this, memory is a basic source of justification. That is a vitally important epistemological point. And as we shall see, the role of memory in our knowledge in general is also of enormous epistemological importance.

Notes

1 We might call merely retained beliefs *weakly* grounded in memory, but I reserve the terms 'memory belief' and 'memorial belief' for beliefs grounded in the normal way illustrated by remembering what I come to believe from, say, perception or testimony.

2 The point that how beliefs are caused, and what their content is, may not indicate how they are grounded (where grounding is the notion crucial to determining whether the belief is justified or represents knowledge) is even wider than so far suggested. A noise too faint for me to hear may cause Tom to jump, which in turn causes me to believe that he is startled; my belief that he is startled is thus (indirectly) caused by the noise, but it is not auditory. It is in no way grounded in my hearing.

3 In both Western philosophy—for example in Plato and Descartes—and Eastern philosophy, innate ideas have played a significant part. In recent times there has been much skepticism about whether they—as opposed to innate dispositions to *form* ideas—are even possible. I cannot discuss this issue here but see no reason not to leave the matter open for the sake of argument. In any case, the possibility of "innate" beliefs seems implicit in something less controversial: that in principle a person could be created as a perfect copy of another, and so would have at least some beliefs at the moment of "birth."

4 The belief is the same in the sense of being an instance of the same cognitive property—the one we both instantiate—as is indicated by our each believing the proposition; but the *instance* of that property is not the same, much as if I raise my hand twice in the same way there are two (act) tokens of the same (act) type.

5 John Locke, in (e.g.) his *Essay Concerning Human Understanding*, speaks of perception as "the inlet of all the materials of" knowledge (Book II, Chapter IX, section 15) and says, comparing perception and memory, that:

> when my eyes are shut or windows fast, I can at pleasure recall to my mind the ideas of light, or the sun, which *former sensation had lodged in my memory* . . . there is a manifest difference between

> the ideas laid up in my memory . . . and those [of perception] which force themselves upon me . . . there is nobody who doth not perceive the difference in himself between contemplating the sun, as he hath the idea of it in his memory, and actually looking upon it: Of which two, his perception is so distinct, that few of his ideas are more distinguishable from one another.
>
> (Book IV, Chapter XI, section 5; italics added)

6 Or virtually the same time: the time-lag argument discussed in Chapter 2 indicates that if light transmission is essential to seeing, there will be a tiny gap between (1) the time at which something we see has a property we are visually caused to believe it has and (2) the time at which we see it as having, or believe it to have, that property. We also found, however, that light transmission does not seem *absolutely* essential for seeing. Note that any causal theory will imply a time-lag when the needed causal connection requires time; but, as with an engine pulling a train, the production of effects by causes need not in all cases require temporal passage.

7 I assume that simple inferences do not require the use of memory; but even if they do, once a belief is *formed* inferentially, it can be inferentially *held* only insofar as it remains supported by the premise beliefs. Then memory may well be what preserves the *inferential structure* represented by believing *p* on the basis of premises; but the belief that *p* is itself only preserved by memory without being genuinely memorial. Not every way that memory preserves a belief renders the belief memorial, and one would explain why one holds this belief not by saying 'I remember . . .' but by citing one's premises.

8 I develop this case, defend the conclusion tentatively stated here, and discuss other matters considered in this chapter in 'Memorial Justification', *Philosophical Topics* 23, 1 (1995), 31–45. For a different position on some of the relevant issues see Carl Ginet, *Knowledge, Perception, and Memory* (Dordrecht: D. Reidel, 1975).

9 It is natural to wonder whether the degree of justification normally belonging to such memory beliefs is as great as that normally belonging to perceptual beliefs. Perhaps not, and one could add 'to some degree' in the normality formulation. But it still appears that the kind of justification is such that it is generally reasonable to believe the propositions in question and that when they are true we commonly can know them on the basis of the relevant kind of justifier.

10 My 'Memorial Justification' (cited in note 8) and some of the literature it refers to consider this difficult question; fortunately, it is not one that requires here any more than the sketch of an answer given.

11 Cf. Michael Huemer's conception of memorial justification in his 'The Problem of Memory Knowledge', *Pacific Philosophical Quarterly* 80

(1999), 346–57 reprinted in his *Epistemology: Contemporary Readings* (London: Routledge, 2002).

12 Note that ascribing justificatory power to memorial images and other memorial experiences, such as the sense of having believed a proposition, *p,* does not commit one to holding that each *time* we have the relevant experience we get *more* justification for *p,* so that we could enhance our justification simply by repeatedly calling up confirmatory images. Granted, a more vivid recollection of the event that *p* represents, such as pruning a tree, may, other things being equal, better justify believing one pruned it than does a less vivid image. We should distinguish *occurrent justification*—the kind based on an experience one is now having—from *dispositional justification*—the kind one has for a belief simply retained in memory. The former varies in degree with the justificatory force of the ground in one's experience at the time; the latter either is invariant, as when one's ground is of a constant justificatory power, or changes only as one acquires new grounds or loses one or more actually possessed. New grounds include new confirmatory evidences; justificatory losses include forgetting or having the confirmatory images one can recall become less vivid. Variability in the force of occurrently justificatory elements does not entail an additive effect. For instance, regarding my belief that I pruned the spruce, calling up my image of pruning it may provide more support on one occasion than on another, perhaps because of differing levels of concentration; but this does not entail that I can parlay my recollections into certainty by simply repeating the exercise, any more than verifying that a pillar remains firmly attached to the corner of one's porch strengthens its support. Indeed, repeating the exercise of recalling an event *can* tend to reduce our justification, at least when it amounts to inflating our self-assurance. I thank Sven Bernecker for raising the difficult question of the justificatory force of repeated recollection.

4 Consciousness

The life of the mind

- **Two basic kinds of mental properties**
- **Introspection and inward vision**
- **Some theories of introspective consciousness**
  Realism about the objects of introspection
  An adverbial view of introspected objects
  The analogy between introspection and ordinary perception
  Introspective beliefs, beliefs about introspectables, and fallibility
- **Consciousness and privileged access**
  Infallibility, omniscience, and privileged access
  Difficulties for the thesis of privileged access
  The possibility of scientific grounds for rejecting privileged access
- **Introspective consciousness as a source of justification and knowledge**
  The range of introspective knowledge and justification
  The defeasibility of introspective justification
  Consciousness as a basic source

4 Consciousness

The life of the mind

So far, we have talked mainly about beliefs concerning things outside ourselves: the green of grass, the smell of roses, the feel of glasses in our hands. But there is much that we believe about what is internal to us. I believe that I am *thinking* about self-knowledge, that I am *imaging* cool blue waters, and that I *believe* I am a conscientious citizen. In holding these beliefs, I attribute rather different sorts of properties to myself: thinking, imaging, and believing. What sorts of properties—or at least phenomena—are they, and how do our beliefs about them give us justification and knowledge? For instance, are some of these self-directed beliefs the products of a kind of *inner perception*? This seems a natural view, and we have already seen how an understanding of perception can clarify memory. If there is some truth in the inner perception view of self-knowledge, exploring the analogy between outer perception and self-consciousness might help to explain how beliefs about our inner lives are justified or constitute knowledge.

Our most important kind of self-knowledge is not about our bodies, but about our minds—for instance about what we believe, want, feel, and take ourselves to remember. It will help to start by describing the kinds of mental properties illustrated by thinking, imaging, and believing. Since they are all broadly mental, this is a task in the philosophy of mind. But epistemology cannot proceed without considerable reflection on mental phenomena, and here it overlaps the philosophy of mind. Thinking, inferring, and believing, for example, are central in both branches of philosophy; and to understand self-knowledge, we need a good sense of what kinds of properties characterize us. We might begin with two kinds that, for our purposes, yield a basic division.

Two basic kinds of mental properties

Thinking is a kind of *process* and involves a sequence of events, events naturally said to be in the mind. Thinking in human beings has a beginning, a middle, and an end; it is constituted by mental events, such as considering a proposition; and these events are always ordered in time, often in subject matter, and sometimes in logic.

Simply *having* an image, in the minimal way one does when there is a static, changeless picture in the mind's eye, is being in a certain (mental) *state*. Unlike something that changes, the existence of such a state does not absolutely require the occurrence of any events during the time it exists. *Imaging* can be a process of calling up a succession of images or, as when one of them is held changeless in the imagination, static. I could image something for a time without any change whatever in my imaging, and without the occurrence of any mental event that might be part of the imaging.

Believing could also be called a mental state, but this terminology can be misleading in suggesting that having a belief is a state of mind, where that implies a global mental condition like worry or excitement. Unlike images and aroused emotions such as jubilation, beliefs do not tend to crowd one another out.

Beliefs differ from images in at least two further ways. First, beliefs need not be in consciousness and indeed we can be conscious of only a quite limited number at once. We all have many which, unlike my belief that I am now writing, we cannot call to mind without some effort. Second, believing need not in any sense be "pictorial." Consider a belief present in consciousness, in the way my belief that the rain has stopped is. This belief is present because I have called it to my attention. I might have held it without attending to it or even to the fact it records.

Even a belief present in consciousness in a prominent way *and* about something as readily picturable as the Statue of Liberty need not involve anything pictorial in the way my imaging must. Suppose I believe that the Statue of Liberty has a majestic beauty standing high in the Bay of New York. Without picturing anything, I can entertain this proposition, and in that way have this belief in my consciousness. By contrast, imaging cool blue waters requires picturing a blue surface. To be sure, when we *call up* this belief about the statue, we tend to picture that structure. But I could later get the proposition in mind, as when I am listing some majestically beautiful landmarks deserving preservation, without picturing anything. I could even retain the belief if I had forgotten what the statue looked like and simply remembered my aesthetic judgment of it.

It will help in sorting things out if we observe a distinction that has already come up but needs more development. Let us call mental properties like beliefs *dispositional* and mental properties like thinking (processes-properties, we might say) *occurrent*. The latter are constituted by mental events and are occurrences: they take place in the way events do and may be said to happen or to go on. The former are not occurrences and may not be said to happen, take place, or go on.[1]

The basic contrast is this. To have a *dispositional property*, or (perhaps not quite equivalently) to be in a dispositional state, is to be disposed—roughly, to tend—to do or undergo something under certain conditions, but not necessarily to be actually doing or undergoing or experiencing something or changing in any way. Thus, my believing that I am a conscientious citizen

is, in part, my being disposed to say that I am one, under conditions that *elicit* that sort of verbal manifestation of my belief, such as your asking me whether I intend to vote. Yet I can have this belief without doing or undergoing anything connected with it, just as sugar can be soluble while it is still in a solid, unaltered lump. I can have the belief in dreamless sleep. By contrast, to have an *occurrent property* is to be doing, undergoing, or experiencing something, as sugar undergoes dissolution. Thus, if you are thinking about mental phenomena you are doing something, even if you are in an armchair; and if you are imaging a flowering crab apple tree, you are experiencing something, at least in the sense that your imaging the tree is now present in your consciousness.

Having a static image, however, as opposed to *calling up* an image, is not a process as, say, silently talking to oneself is. Occurrent mental properties, then, must be subdivided. To mark the difference, we might call occurrent mental properties like thinking *experiential process properties* and occurrent mental properties like having a static image in mind *experiential state properties*.[2] Clearly, both differ from dispositional mental properties; possessing those does not even require being conscious, much less having a kind of experience. All three kinds of mental properties turn out to be important for understanding the epistemological role of introspection.[3]

Introspection and inward vision

If we take a cue from the etymology of 'introspection', which derives from the Latin *introspicere*, meaning 'to look within', we might construe introspection as attending to one's own consciousness and, when one's mind is not blank, thereby achieving a kind of inner seeing. I might introspect my images, for instance, and conclude that my image of the spruce and nearby maple indicates that the spruce is taller than the maple. I might *have* to introspect my image of the maple to tell without looking back at it whether it has three secondary trunks. Introspection need not, however, be labored or even constitute an act. It may be simply a matter of becoming conscious of something in my mind. This can be as natural as something's coming into one's physical field of vision, rather than like making the effort of observation in order to see.

It is not only in consciously introspecting that one can vividly image. In Shakespeare's *King Lear* there is a scene in which Edgar wants to convince Gloucester, who has lost his sight, that he is at the top of a cliff. Edgar's description is so vivid that the deception succeeds:

> How fearful and dizzy 'tis to cast one's eye so low!
> The crows and choughs that wing the midway air
> Show scarce so gross as beetles. Halfway down
> Hangs one that gathers samphire, dreadful trade!
> Methinks he seems no bigger than his head.

> The fishermen that walk upon the beach
> Appear like mice, and yond tall anchoring bark . . .
> Almost too small for sight.
>
> (Act IV, scene vi)

What Gloucester sees in his mind's eye is so vivid that he believes he is at the edge of a precipice. His visual consciousness is filled with images from Edgar's portrait. Here introspection is simply a matter of vivid consciousness of the imagery that is before the mind.

If introspective consciousness does produce inner seeing and other sensuous imagery (such as, commonly, sound), we can try to understand it by drawing on what we know about perception. For instance, we can explore introspectional counterparts of some theories of perception and sensory experience. But one limitation of that procedure is apparent the moment we reflect on the dispositional mental properties, for instance believing, wanting, or having a fear of cancer. We do not *see* such properties in any sensory way, nor even as we may be thought to see (in our mind's eye) an image of cool blue waters.

The analogy to vision might, however, still hold for introspection regarding *occurrent* mental properties. If it does, it presumably applies only to the mental state properties, such as imaging. For surely thinking is not seen. It need not even be heard in the mind's ear. I may hear my silent recitation of Shelley's 'Ozymandias', but thinking *need* not occur in inner speech, certainly not speech of that narrative, punctuated sort.

Perhaps it is only pictorial mental properties that we see through inner vision; and perhaps it is only other "phenomenal" properties, such as inner recitations, tactual imagings (say, of the coldness of a glass), and the like, that seem accessible to inner analogues of perception: to hearing in the mind's ear, touching in the tactual imagination, and so on. It is doubtful, then, that we can go very far by conceiving introspection as simply producing inward seeing. Still, it is worth exploring how the analogy to seeing holds up for the one important case of pictorial properties.

Some theories of introspective consciousness

Suppose that introspecting such things as images of cool blue waters does produce a kind of inner seeing. Are we to understand this seeing on realist lines, so that there must be some real object seen by the introspective eye?

Realism about the objects of introspection

One might think that the sense-datum view simply cannot be extended in this way to introspection. This is at least a natural assumption about self-understanding. For on the introspectional counterpart of the sense-datum view, seeing (in one's mind's eye) an image of cool blue waters would require

something like *another* image, one that represents the first one in the way sense-data represent a physical object seen by virtue of the perceiver's acquaintance with them. Call it a *second-order image*, as it is an image of an image.

What would second-order images be like? If I try to image my image of cool blue waters, I get that very image all over again, or I image something else, or I get something that is not an image at all, such as a *thought* of my original image. But this point does not show that there could not be second-order images. Perhaps there could be some that are less vivid than the originals they picture, just as my imaginational image of blue waters is less vivid than the sensory image I have in seeing those waters. In any case, sense-datum theorists could avoid positing second-order images. It would be more plausible to hold that to image simply *is* to be acquainted with sense-data of a kind different and typically less vivid than outer perception yields.

An adverbial view of introspected objects

A defender of an adverbial account of sensory experience would not countenance images as sense-datum-like entities with properties in their own right. Take first perceptual imaging that is later "copied" in retrospective imagination. Adverbialists will likely hold that there is really just *one* basic kind of imaging process, and that it occurs more vividly in perception than in imagination. Thus, imaging blue waters is simply imaginationally, rather than perceptually, sensing in the way one does upon seeing blue waters: sensing "blue-waterly," as we might adverbially express it. Since the adverbial view conceives imaging as a way of experiencing rather than as a relation to an object, there *is* no image as an object to be copied.

On the adverbial view, then, there is no need to posit second-order images to represent first-order mental images to us, and the less vivid imagings which might seem to represent mental images are best construed as less vivid occurrences of the original imaging process. This point does not show that there *cannot* be second-order images or similar interior objects of inner vision, but the adverbial view reduces the inclination to think that there are any by suggesting a plausible alternative account of the facts needing explanation. Chief among these facts is that in recalling an image (especially a sensory image), one may have a less vivid image which apparently stands to the former as an imaginational image of a scene stands to the sensory image of that scene from which the imaginational image seems copied. The adverbial account of sensory (and other) experiences might explain this by interpreting the recalled image, say of blue waters, as *recollectively* sensing blue-waterly, where this is like visually sensing blue-waterly, but less vivid.

Given these and other points, it seems doubtful whether any realist theory of the introspection of images—one that takes them to be objects existing in their own right and having their own properties—can justify a strong analogy between that kind of introspection and ordinary viewing. For it is

by no means clear that there is any *object* introspected to serve as the counterpart of an object of ordinary vision. For an adverbial approach to experience, although realism about the (physical) objects of perception is a highly plausible view, realism about the objects of introspection is not. The idea is roughly that mental properties, such as imaging, can adequately represent physical objects in our mental life; inner objects should not be postulated for this task.

The anti-realist element in this adverbial view should not be exaggerated. To deny that mental images are objects having their own properties and in that sense are not real does not in the least imply that *imaging* is not real. Imaging processes are surely real properties of persons, even though they are not relations between persons and objects of immediate, inner perception. Does this imply, say, that introspection has no object in the sense of something it is *of* (or about), such as imaging blue waters? Certainly not. But this "object" is a kind of *content*, not an entity to which the mind is related. On the adverbial view of introspection, this kind of object is determined by what the introspection is about and is not a thing with its own properties, such as colors and shapes, sounds and movements, depths and textures.[4]

The analogy between introspection and ordinary perception

The adverbial view in question may seem unable to do justice to the apparently causal character of introspection. There is surely some causal explanation of my being acquainted with, say, imaging blue waters rather than imaging the Statue of Liberty when I monitor a daydream of a rural summer holiday. Perhaps it is mainly in what causes the relevant imaging that such introspective consciousness differs from seeing. How might this difference be explained?

Suppose the adverbial account is true. Introspection may still be like simple perception in two ways. First, introspective viewing may imply some kind of causal relation between what is introspected in it, say an imaging, and the introspective consciousness of that state or process. Second, such viewing may imply a causal relation between the object of introspective knowledge—for instance one's imaging blue waters—and the beliefs constituting this knowledge.

In explaining the analogy between introspection and perception, I want to concentrate mainly on introspective beliefs as compared with perceptual beliefs; we can then understand how introspection, and indeed consciousness in general, can ground justification and knowledge. A major question is how we can determine whether what the theory says is true: whether, in introspecting, as when we concentrate on our own imaging, the beliefs we thereby form about what we are concentrating on are produced by that very thing or by some aspect of it such as its imagined blue color. Only to the extent that they are should we expect introspection to ground justification and

knowledge in the broadly causal way that perception does. Many considerations are relevant here, but let me cite just two sorts.

First of all, it is surely *because* I am imaging cool blue waters that, when, with closed eyes, I introspectively consider what I am conscious of, I believe I am imaging them (and am conscious of my imaging them). It is natural (and reasonable) to take this 'because' to express a causal relation, even if I *could* be mistaken in thinking there is one. If the causal basis of my belief is not some inner object seen (as on the sense-datum theory), it is presumably the state or process of imaging. This is, in any event, how an adverbial theory of sensory experience would view the causal relations here. Similarly, if I introspectively believe that I am thinking about introspection, I believe this because I *am* thinking about it: it is the thinking itself that causes me to believe that it is occurring. In both cases, the introspective beliefs are produced by inner processes, and indeed in a way that makes it plausible to regard the beliefs as true. Some inner processes are *like* seeing an object in still other ways, but these processes can all be understood without positing inner objects analogous to perceptible ones such as trees and seen by the introspective eye.

A second point is this. Suppose my believing that I am imaging cool blue waters is not caused by my imaging them. The belief is then *not* introspective at all. It is *about* what is introspectable, but it is not *grounded* in introspection, any more than a belief merely about a perceptible, such as the rich red in a painting in a faraway museum, is a perceptual belief. Here, then, is another important similarity between introspection and perception.

Introspective beliefs, beliefs about introspectables, and fallibility

It may seem that the case described—believing one is imaging something when in fact one is not—is impossible. But suppose I have been asked to image cool blue waters, yet I hate the water and anyway have a lot on my mind. Still, if I want to be cooperative, then even though my mind is mainly on my problems, I may call up an image. However, as I am not concentrating on calling up the image, the image that I actually get might be only of a blue surface, not of blue waters. I might now inattentively assume (and thereby come to believe) that I have called up the requested image of cool blue waters. This belief is produced by a combination of my calling up the wrong image, which I do not attentively introspect at all, and by non-imaginational factors such as my desire to cooperate. I might even retain the belief for at least some moments after I cease to image at all. In that case, it is not only not true; it is not even introspective.

This example suggests that even a true belief about one's conscious states or processes would not be introspective without being causally connected with them. It would be about these introspectable elements but not grounded in "seeing" them in the way required for being an introspective belief. Other examples support the same point. Imagine that my task is to think about

introspection for a solid hour. I monitor myself and, on the basis of introspection, conclude from time to time that I am thinking about introspection. As I reflect on my topic, I continue to believe that I am thinking about introspection. Now when I truly believe this simply because I have repeatedly confirmed it and am confident of steady concentration, and *not* because I am still monitoring myself introspectively, my belief, though perfectly true, is not introspective.

The best explanation of this point seems, again, to be that my belief is not caused (in the right way, at least) by the thinking that should be its ground. It is a retained belief about my ongoing mental activity; it is not *grounded* in that activity as a focus of my introspective attention. My belief that I am thinking about introspection is a propositional belief *that* I am now doing so, but it is not an objectual belief, regarding my present thinking, to the effect that it is about introspection. It is not grounded in my *present* thinking at all, any more than my belief that a painting I remember portrays a rich red coat is grounded in seeing that red.

An overall conclusion we may draw here is that although there may be no objects such as sense-data or imaginational copies of them which we introspect, the process by which self-consciousness leads to introspective beliefs, and thereby to knowledge and justified beliefs about one's own mind, is nevertheless causal. Like perception of the outside world and (though in a different way) recalling events of the past, it produces something like a sensory impression and, at least commonly, beliefs about what seems to be revealed to one by that impression. The causes of introspective beliefs, however, are apparently processes and events in the mind. They are not, or at least need not be, objects that reside therein.

Consciousness and privileged access

Suppose that introspective consciousness is causally grounded in the way we have seen. We should then raise some of the same epistemological questions about it that we raised about perception. For instance, is introspection subject to counterparts of illusion and hallucination? And if it is, how might it still be a source of justification and knowledge? Let us start with the question of how anything like illusion or hallucination might occur in consciousness.

Infallibility, omniscience, and privileged access

One might think that, regarding the inner domain, which is the subject of introspective beliefs, one cannot make mistakes. If so, one might conclude that illusion and hallucination regarding this domain are impossible. Indeed, David Hume maintained that since "the contents of the mind" are known by "consciousness" (by which he meant something at least much like introspection), "they must appear in every respect what they are, and be what they appear."[5]

Hume's statement suggests two far-reaching claims about self-knowledge. One claim—that the contents of our minds must be what they appear to us to be—expresses the idea that introspective consciousness can give us beliefs that cannot be mistaken. The other claim—that these contents must appear to be what they are—expresses the idea that consciousness makes us so richly aware of the (introspectable) contents of the mind that it guarantees us full knowledge of them. These ideas need refinement before we can reasonably appraise them.

The first claim suggests a *thesis of infallibility* (impossibility of error): an introspectional belief—roughly one to the effect that one is (now) in an occurrent mental state (such as imaging) or that one is undergoing a mental process (such as thinking) or that one is experiencing something (such as pain)—cannot be mistaken. The infallibility thesis rests largely on the idea that we are in such a strong position regarding occurrent mental phenomena that we cannot err in thinking they are going on inside us.

The second claim suggests a *thesis of omniscience* (all-knowingness) with respect to the present occurrent contents of consciousness: if one is in an occurrent mental state, undergoing a mental process, or experiencing something, one cannot fail to know that one is. The omniscience thesis rests largely on the idea that occurrent mental phenomena are so prominent in consciousness that one cannot help knowing that they are present.

Together, these two theses constitute the *strong doctrine of privileged access*. The infallibility thesis says that our access to what is (mentally) occurring in us is so good that our beliefs about its present make-up are infallible; there is no risk of *error*. The omniscience thesis says that our access to it is so good that we cannot fail to know what (mentally) occurs in us; there is no risk of *ignorance*. It is because no one else is in such a good position to know about our mental life, and because we ourselves are not in such a good position to know about the *external* world, that it is natural to speak here of *privileged* access. The strong doctrine of privileged access is associated not only with Hume but even more with René Descartes, who is widely taken to maintain it in his famous *Meditations on First Philosophy* (1641), especially in Meditation Two.

Suppose for the sake of argument that the infallibility and omniscience theses are true. One might think that this would preclude inward counterparts of illusion and hallucination. It would not. Having illusions and hallucinations does *not* imply having false beliefs or ignorance of anything. Looking from a sharp angle from corner to corner, you can see a book as having the shape of a (non-rectangular) parallelogram without believing that it has that shape; and I can hallucinate a spruce tree like one that has burned to the ground without believing it is before me. We may know the facts. For inner perception as for ordinary sense perception, phenomenal experience is one thing and belief another.

Suppose, on the other hand, that there *are* no inner objects, such as blue, watery images, to appear to us to have properties they do not possess, such as

wavy surfaces. Then illusions of the kind we have in perception, in which an object appears to have properties it actually lacks, could not occur, as there is no object to appear to us. Nor could a hallucination of, say, an image of blue waters be *of* such an object and true or false *to* it, though of course in a sense it "misrepresents" reality. Suppose, however, that there are inner objects that we see when we image. How would hallucinating an image of, for instance, a loved one differ from just *having* that image? A sense-datum theorist might say that the hallucinatory image would be less vivid or stable than a real one. But it would still be an image of the same thing and could also be just like a normal image in vividness and other respects. It would be wrong to say, then, that a hallucinatory image is necessarily a less vivid or less stable version of a normal image, and the difficulty of explaining the difference between hallucinatory and real images is an additional reason to avoid (as the adverbial view does) positing mental images as objects.[6]

Difficulties for the thesis of privileged access

It might be, however, that quite apart from illusion or hallucination, we can have false beliefs, or suffer some degree of ignorance, about our mental life. This is clear for *some* mental phenomena, such as dispositions like believing, wanting, and fearing. We can mistakenly believe that we do not have a certain ignoble desire (say to make a fool of a pretentious boss), particularly if it is important to our self-image that we see ourselves as having no hostile desires. For the same reasons, we can fail to know that we *do* have the desire. One can also discover a fear which, previously, one quite honestly disavowed because it was at odds with one's sense of oneself as courageous.[7]

Dispositions, however, should not be conceived as *occurring*, and in any case it is the occurrent mental phenomena to which philosophers have tended to think we have the kind of privileged access expressed in the theses of infallibility and omniscience. Can we be mistaken, or at least ignorant, about our occurrent mental states or processes?

Consider first the possibility of mistake. Could I believe I am thinking about the nature of introspection when I am only daydreaming about the images and feelings I might introspect? It seems so, provided I do not attend closely to what is occurring within me. Granted, this would not be a wholesale mistake; it would be like thinking I am watching someone observing a game, when I have become preoccupied with the game itself and have ceased to attend to its observer.

Suppose, however, that the infallibility thesis is restricted to beliefs based on *attentive* introspection, where this implies "looking" closely at the relevant aspect of one's consciousness. Call this the *restricted infallibility view*; it says only that *attentive introspectional beliefs are true*. If I carefully consider the proposition that I am thinking about introspection, and I believe it on the basis of attentive introspection (i.e. on the basis of my carefully focusing on the relevant aspect of my consciousness), could this belief be mistaken?

It may seem that error here is impossible. But suppose I desperately want to believe that I am thinking about introspection. Could this not lead me to take my daydreaming to be such thinking and even cause me to form an attentive introspective belief that I am doing such thinking? It seems so. Similarly, I could believe, on the basis of attentive but imperfect introspection, that I am imaging an octagon and then, concentrating harder and counting sides, discover that the figure has only seven.

If, for some occurrent mental states (such as thinking), it is possible to be mistaken in believing that one is now in them, then the omniscience thesis of privileged access should be abandoned along with the infallibility view. This holds even if the omniscience thesis, too, is restricted, as it should be, to cases of carefully attending to one's consciousness. The easiest way to see why fallibility cuts against omniscience is to note how omniscience would tend to guarantee *in*fallibility and so would be cast in doubt if the latter is. Let me explain.

Given the extensive self-knowledge implied by omniscience, if, instead of thinking about the nature of introspection, I am only daydreaming, then I must know that I am daydreaming. But I will presumably not be so foolish as *also* to believe that I am thinking about introspection—something plainly different from daydreaming. Since I would know as well that I am occupied with, say, a series of images that portray me as swimming in cool blue waters, it is even less likely that I will believe I am thinking about introspection. It appears, then, that if I know every truth about—am omniscient about—my consciousness, then I presumably cannot believe any falsehood about it and so am infallible about it as well.[8]

It is at best extremely unlikely (and perhaps impossible) that these two things—knowing every truth about one's consciousness and nonetheless believing some falsehood about it—occur together, leaving one omniscient regarding one's own consciousness, yet inconsistent and fallible about it. One would know every truth about it yet would also somehow believe falsehoods incompatible with those truths. This being at best improbable, if I am fallible I am at least very likely not omniscient. Now recall our daydreaming example. It casts doubt even on the restricted thesis of omniscience. In that example, although I am in fact daydreaming, I would presumably not know that I am. If I do know that I am daydreaming, I would believe this, and then it is very doubtful that I would *also* believe I am thinking about introspection.

These points suggest that, contrary to the thesis of omniscience, I can fail to know certain things about my consciousness even when I am attending to it; but they do not imply that the omniscience side of the privileged access view is wildly mistaken, in that I might be ignorant of *every* truth about my daydreaming. Far from it. As I (objectually) believe the daydreaming to be thinking about introspection, I presumably at least know my daydreaming to involve words or colors or shapes. I have some knowledge *of* it, but I would still not know the proposition *that* I am daydreaming and thus would not be omniscient regarding the mental processes occurring in me.

The possibility of scientific grounds for rejecting privileged access

Perhaps there could someday be a source of significant evidence against even the restricted doctrines of privileged access. For it could turn out that every occurrent mental phenomenon is uniquely correlated with some distinct set of brain processes. Then someone could devise a "cerebroscope" for viewing the brain and could read off the contents of consciousness from the cerebroscopic data (a possibility with disturbing implications that require ethical scrutiny). What would guarantee that our introspective beliefs must match what the machine says about our mental lives? And what would a mismatch show?

Imagine that we could discover cerebroscopically a unique neural pattern for, say, believing on the basis of attentive introspection that one is imaging cool blue waters, at the same time as we discover the pattern for imaging a field of blue-green grass. It would be natural here to suppose the subject is mistaking the grassy imaging for a watery one. Might we not regard the sophisticated equipment as more likely to be right than the subject?

There is a problem with this reasoning. How could one *establish* the unique correlations except by relying on the accuracy of people's introspective beliefs? Might it not be necessary to start by *asking* people what they are, say, imaging, to assume that they are correct, and only *then* record the associated brain state? And if learning the correlations would depend on the accuracy of introspective reports, how could the correlations show such reports to be mistaken?

A possible reply is this. First, let us suppose that learning the correlations would depend on the accuracy of introspective reports. Still, neuroscientists would not have had to rely on the accuracy of precisely the introspective belief being shown to be mistaken, and perhaps not even on the accuracy of highly similar beliefs. In any event, once they construct their instrument, they might no longer need to consult introspection to use it. They might throw away the very ladder they have climbed up on.

Imagine, however, that they do have to rely on just the sorts of belief we are examining, together with evidence regarding these beliefs' reliability—evidence we already have independently of the cerebroscope. Would this imply that the cerebroscope could not provide powerful evidence against introspective beliefs?

Consider an analogy. We might use a mercury thermometer to construct a gas thermometer. We might calibrate a container of gas with a piston that rises and falls as the gas is heated and cooled. The new temperature readings might correlate perfectly with mercury readings in many instances: in measuring water temperature, wood temperature, and in other cases. The gas thermometer might then be used for the same jobs as the mercury thermometer *and* might gauge temperatures that the mercury thermometer cannot measure, say because they are above its boiling point. Could we not use a gas

thermometer to correct a mercury thermometer in some cases, or perhaps to correct all mercury thermometers within limits? We could. This seems so even if we originally thought the mercury thermometer infallible in measuring temperature, perhaps because we mistakenly considered its readings partly definitive of what temperature *is*. We would rebuild the ladder we have climbed up on.

Similar points might hold for beliefs about what is now occurring in one. If the analogy does extend this far—if the gas thermometer is to the mercury thermometer rather as the cerebroscope is to sincere testimony about current mental states—then even the restricted omniscience view fares no better than the restricted infallibility view. For even when I am attentive to what is in my consciousnes, a cerebroscope could indicate that I do not believe (hence do not know) that a certain thing is occurring, such as a frightening image which I believe I have put out of mind.

Introspective consciousness as a source of justification and knowledge

It is important not to overextend our criticism of various claims of privileged access. After all, even the restricted infallibility and omniscience views are very strong claims of privileged access. They can be given up along with the strong theses of privileged access quite consistently with holding that our access to what is occurring in us is very privileged indeed. Let us explore this.

The range of introspective knowledge and justification

Nothing I have said undermines a qualified epistemic principle. This *self-knowledge principle* says: our attentively formed introspective beliefs about what is now occurring in us are *normally* true and constitute knowledge. The difficulty of finding grounds for thinking they even *could* be false provides some reason to consider them at least very likely to be correct. Similarly, when we are attentive to what is occurring in us, then if something (knowable) is occurring in us, such as a certain melody in the mind's ear, *normally* we know that it is occurring, or at least we are in a position to know this simply by attentively forming a belief about what is going through our mind. At least this qualified epistemic principle holds for the domain of our conscious life.

Granted, our "access" to our dispositional properties is not as good as our access to what is occurring in us. We need not be conscious of the former properties, whereas the very existence of one's imaging (or of an image if there are such objects) *consists in* its place in consciousness. Beliefs and other mental dispositions need not even enter consciousness, or ever be a subject of our thoughts or concerns. Some of them may indeed be *repressed*, so that we normally cannot easily become aware of them.[9]

Nevertheless—and here is a *justification principle applicable to the*

dispositional mental domain—our beliefs to the effect that we are now in a dispositional mental state, for instance wanting, fearing, intending, or believing, are normally justified. We might also say that such beliefs, though defeasibly justified, are always prima facie justified, so that they are justified overall unless some defeating factor, such as an abnormal psychological interference, occurs. Moreover, normally, when we have a want (or fear, intention, belief, or similar disposition), we are in a position to know (and justifiedly believe) this. We can, then, usually know this if we need to. We very commonly do not know it, however; for such things may not enter consciousness at all, and there is often no reason to take any notice of them or form any beliefs about them. This kind of ignorance is innocuous.

There are many issues and details I have not mentioned; but we can now generalize about introspection (roughly, self-consciousness, i.e., consciousness turned inward) in relation to belief, justification, and knowledge, and summarize our main epistemological conclusions. Plainly, many beliefs arise from introspection, and the points that have emerged suggest a *general epistemic principle concerning self-knowledge* which, though far weaker than the infallibility thesis, is far-reaching: *normally, introspective beliefs grounded in attentive self-consciousness are true and constitute knowledge.* (This principle is slightly different from the self-knowledge one stated above.)

A second epistemic principle—*an attentional epistemic principle concerning self-knowledge*—though far weaker than the omniscience thesis, is that normally, if we attentively focus introspectively on something going on in us, we know that it is going on, under at least some description. I may not know that I am humming the slow movement of Beethoven's Pathétique Sonata, but I do know I am humming a melodic piano piece.

The corresponding justification principles suggested by our discussion seem at least equally plausible. First—to cite an *introspective justification principle*—normally, introspective beliefs grounded in attentive introspection are justified; and, second, normally, if I attentively focus on something going on in me, I am justified in believing that it is going on in me. To be sure, some such beliefs are better justified than others, and even some that are not attentive are justified. All of them are plausibly regarded as prima facie justified.

There are many possible principles regarding our justification and knowledge about ourselves, and there are many possible qualifications of the one just stated. But those principles are sufficient to suggest the power of introspection as a source of justification and knowledge. The examples I used to argue that introspection is fallible do not show that the apparently false introspective beliefs were *unjustified* or that true ones are not knowledge. A false belief, particularly if it is of a kind usually justified, can still be justified; and a true belief of a kind that can sometimes be false may itself constitute knowledge.[10]

The defeasibility of introspective justification

These points about the high degree of privileged access we apparently do have may create a danger of overestimating the strength of introspective justification. From our examples, it might be thought that attentive introspection, even if not absolutely infallible, generates a kind of justification that at least cannot be defeated. Even if I am somehow mistaken about whether I am imaging blue waters, if I believe this on the basis of introspection, it would seem that I am *in the right*, even if objectively I am *not right*.

How could I fail to be justified in believing that I am imaging cool blue waters, if my belief is grounded in attentive introspection? If the question seems rhetorical, this may be because one thinks there simply is nothing else I should have done besides attending and hence no possible defeaters of my justification by appeal to another kind of ground for belief. Let us explore this.

Granting that I could not fail to be justified *unless* I could have good reason to believe I may be mistaken, still, perhaps I could fail to be justified if I had good reason for believing I am mistaken, such as evidence from repeated cerebroscopic results indicating that I have been mistaken in many quite similar cases. It is far from obvious that I could not have sufficient evidence of this sort. It seems wisest, then, to conclude that although introspective justification tends to be very strong, it remains prima facie rather than absolute and can be defeated by counter-evidence.

In any case, plainly beliefs grounded in attentive introspection, such as my belief that I am now imaging blue waters, are normally justified to a very high degree. Moreover—and here we have still another justification principle—*normally, my simply being engaged in attentive introspection also yields situational (propositional) justification for beliefs about what I am attending to*, even when it does not in fact yield any such beliefs. If I somehow "notice" my imaging blue waters yet do not form the belief that I am doing so, I am nonetheless (prima facie) justified *in* believing, and have justification for believing, that I am, just as, if I see a bird fly past and take no notice of it, I am still justified in believing it is flying past me. The analogy to perception seems sound here, and that is one reason why introspection is considered a kind of inner observation and (unless it somehow yields no content) a kind of inner perception.

Consciousness as a basic source

If we now ask whether consciousness, including especially introspective consciousness, is a *basic* source of belief, justification, and knowledge, the answer should be evident. It is. In this, as in many other respects, it is like perception. But it may well be that the degree of justification which consciousness (including introspection) generates is greater than the degree generated by perceptual experience, other things being equal.

The special strength of justification on the part of beliefs about elements in consciousness has led some philosophers to think that these beliefs are a kind of foundation for knowledge and for the justification of all other beliefs. Descartes is often thought to have so regarded introspectively grounded beliefs or knowledge. Whether knowledge and justification need a kind of foundation and whether, if they do, these beliefs are the best candidates to serve as a foundation—better than, say, perceptual and memory beliefs—are the major questions pursued in Chapter 9.

There is a further epistemologically significant difference between perception and consciousness, especially as manifested in introspection, as sources of knowledge (and justification). We can by and large introspect *at will*—roughly, just by (sufficiently) wanting to—though we may also do it quite spontaneously. Moreover, there is no limit to how many things we can come to know by introspecting, if only because we can, without limit, call up images and construct thoughts. But we cannot perceive at will; and what we can know through perception is limited by what there is outside us to perceive and by external conditions of observation, just as what we can know through remembering or recalling is limited by what has actually happened (or what propositions are true) and by the conditions of belief or image retention crucial for remembering or recalling.[11]

Introspective consciousness, then, is unlike perception and memory in enabling us to acquire a considerable amount of knowledge whether external circumstances cooperate or not. Whatever one can "observe" in one's own mind is a possible subject of study, and it appears that many of the beliefs we attentively form concerning our mental lives tend to constitute genuine knowledge. Very roughly, introspective consciousness is a substantially *active faculty*; perception and memory are largely *reactive faculties*.

Granted, some content—like sensations of pain—comes into consciousness uninvited. Still, we can very freely *call to mind* both propositional and imagistic content. Some of it may come from memory, which shows that introspective consciousness may draw on that as well on resources created by imagination or intellect. By contrast, sensory content, such as perceptual images, enters our mind only when our senses are *taken*, by our own observational efforts or by contingencies of experience, to it. In the inner world, by sharp contrast with the external world, there is far more at our beck and call. This is perhaps another reason why introspectively grounded beliefs have sometimes seemed to be such good material to serve as foundations for knowledge and justification. In addition to the high degree of justification self-consciousness commonly confers on beliefs, it is an active source of both justification and knowledge.

There is a trade-off, however. Through perception, we acquire (primarily) justified beliefs and knowledge about the external world; without these, we would likely not survive. Through introspection, we acquire (primarily) justified beliefs and knowledge only about the internal world; with only this, our knowledge and justification would be sadly limited, even if we could survive.

This is not to underplay the importance of the internal world: without good access to it we would have little if any self-knowledge and, for that reason, probably at best shallow knowledge of others.

Self-knowledge is also important as a back-up when questions arise about one's justification or knowledge regarding external objects. Confronted with a strange object, one may carefully consider the stability, coherence, and variations of one's perceptual experiences of it in order to rule out hallucination. Told that one merely imagined a car's passing, one may try to recall the event and then scrutinize both the vividness of one's imagery and one's confidence that the belief comes from memory rather than merely imagination. Without the kind of self-knowledge possible here, we would have less knowledge about the external world. Both perceptual and introspective knowledge are vital, and both, as we shall soon see, can be extended, by good reasoning from the raw materials they supply, far beyond their beginnings in our experience.

Notes

1 To say that beliefs are dispositional properties does not imply that the concept of believing is dispositionally *analyzable,* hence equivalent to a set of conditional propositions; and I do not think the concept is dispositionally analyzable.
2 To be sure, images can be possessed memorially, as is my image of the Statue of Liberty when I do not have it in mind; and 'imaging' can designate a process, as when I call up the series of images corresponding to looking at the statue from the Brooklyn Heights Promenade and glancing northward to Lower Manhattan, thence to the Brooklyn Bridge, and up the East River beyond the bridge.
3 Both kinds of properties are experiential, in that they represent features of experience. Both, then, might be considered phenomenal, but sometimes the term 'phenomenal property' is restricted to the sensory kind that characterizes either the five senses or "inner sense," by which we feel sensations of pain and pleasure.
4 Such contentual objects are often called *intentional objects*, largely on the ground that, like lofty deeds we intend to perform but do not do, they need not exist.
5 See David Hume, *A Treatise of Human Nature* (first published in 1739–40) (Part IV, section II), ed. by L.A. Selby-Bigge (Oxford: Oxford University Press, 1888).
6 One might still distinguish genuine from hallucinatory images by insisting that to be a genuine image *of* (say) a loved one is to be an image caused by the corresponding sense, say seeing that person. But this view has an odd consequence. Through hearing a detailed description I could have an accurate image of Maj that is in a sense *of* her, as it perfectly "pictures" her, even if I have never seen her; but this would be mistakenly classified

as a hallucinatory image, by the causal conception just stated. There are certainly different kinds of images and various ways in which they can mislead, but the analogy between perception and introspective consciousness does not extend in any simple way to the possibility of inner illusions and hallucinations. I cannot pursue the matter further here, but for a detailed non-technical discussion of mental imagery see Alastair Hannay, *Mental Images: A Defence* (London: George Allen & Unwin, 1971) and my critical examination of this book in 'The Ontological Status of Mental Images', *Inquiry* 21 (1978), 348–61.

7 Some of these cases seem to occur in *self-deception*, a phenomenon that raises profound questions for both epistemology and the philosophy of mind. For a comprehensive collection of papers on it (including one offering my own account), see Brian P. McLaughlin and Amélie O. Rorty (eds.), *Perspectives on Self-Deception* (Berkeley and Los Angeles, CA: University of California Press, 1988).

8 The thesis of omniscience might be restricted to *introspectable* truths, as opposed to such truths as that there are 101 berries visible on the blackberry bush I am imaging, which I could know only on the basis of memory (and arithmetic) as well as introspection. The infallibility thesis might also be plausibly restricted in a similar way. This point bears on the connection between the two theses but should not affect the argumentation in the text.

9 Repression need not be exactly the kind of thing Sigmund Freud described, requiring psychoanalysis or special techniques to come to consciousness. There are various kinds and degrees of repression; the point here is simply that *having* a belief (or other dispositional state) is possible even if it is repressed. One might, for example, still act in the way expected of a believer of the relevant proposition. Such action is one route to discovering repressed beliefs.

10 For reasons to be considered in Chapter 13, skeptics tend to deny this.

11 There is less disanalogy in the negative cases: we cannot always cease at will to concentrate introspectively on our mental life, as illustrated by preoccupying pains; and we cannot in general cease perceiving at will; we must, for example, do so by closing our eyes or turning off a radio. This blocks the path of observation, just as an aspirin might block the path of pain.

5 Reason I

Understanding, insight, and intellectual power

- **Self-evident truths of reason**
  The concept of self-evidence
  Two types of immediacy

- **The classical view of the truths of reason**
  Analytic propositions
  Necessary propositions
  The analytic, the a priori, and the synthetic
  Three types of a priori propositions
  The empirical
  Analytic truth, concept acquisition, and necessity

- **The empiricist view of the truths of reason**
  Rationalism and empiricism
  Empiricism and the genesis and confirmation of arithmetic beliefs
  Empiricism and logical and analytic truths

5 Reason I

Understanding, insight, and intellectual power

I see the green field and I believe that it is there before me. I look away, and I believe that I am now imaging it. I remember its shape, and I believe that it is rectangular. These are beliefs grounded in my experience: perceptual, self-conscious, and memorial. But I also believe something quite different: that *if* the spruce to my left is taller than the maple to my right, then the maple is shorter than the spruce.

On what basis does one believe this obvious truth? Do we even need to see the trees to know it? Certainly it is on the basis of perception that I believe *each* of the two comparative propositions; it is easy to see, for instance, that the spruce is taller than the maple. But I do not believe on the basis of perception that *if* the spruce is taller than the maple then the maple is shorter than the spruce. As a rational being, I quite easily grasp this truth and thereby believe it.

The kind of apparently elementary use of reason this case illustrates seems basic for both knowledge and justification. But there is continuing debate about the nature and grounds of our knowledge and justification regarding the simple, obvious truths that we seem to know just in virtue of the kind of understanding of them any rational being might be expected to have. A good way to understand the epistemological role of reason is to begin with a notion that seems central for the most basic kind of knowledge and justification reason gives us—*self-evidence*.

Self-evident truths of reason

Such truths as the luminous one that if the spruce is taller than the maple then the maple is shorter than the spruce have been said to be evident to reason, conceived roughly as a mental capacity of understanding. They are presumably called *self-evident* because they are thought to be evidently true taken by themselves, with no need for supporting evidence. Indeed, they are often considered *obvious* in themselves, roughly in the sense that simply upon attentively coming to understand them, one normally sees their truth and thereby knows them.

The concept of self-evidence

In the light of such points, we might more specifically characterize self-evident propositions as those truths such that (1) if one (adequately) understands them, then by virtue of that understanding one is justified in (hence has justification for) believing them, and (2) if one believes them on the basis of (adequately) understanding them, then one thereby knows them.[1] (1) says roughly that understanding them suffices for being situationally justified in believing them; it provides a justification for belief—which one can have without actually believing the proposition in question. (2) says in effect that this understanding can ground knowledge: the understanding is sufficient to render a belief based on it knowledge. (2) implies, then, that self-evident propositions are true. This implication is appropriate, since the self-evident is standardly regarded as true (and for clarity I have put truth explicitly into the characterization above).

What I have said does not imply, however, that the kind of justification one gains from understanding the self-evident is *indefeasible*, that is, so secure that it cannot be defeated, rather than prima facie. If the understanding in question is eliminated or obscured, the belief may cease to be justified. But at least some cases of this kind of justification are plausibly taken to exhibit justification as strong as any we can have. It can be difficult to see how defeasibility can occur here because self-evident truths are so commonly considered *also* obvious. But not all of them are—at least to finite minds. Apart from logical training, certain self-evident logical truths are not obvious to us; and it may not be obvious to most of us, on first considering it, that first cousins share a pair of grandparents. But this satisfies both (1) and (2) and is self-evident.[2]

There is an important analogy to perception. Just as one can see a visible property of something, such as its rectangularity, without believing that it has that property, one can comprehendingly (understandingly) consider a self-evident proposition without coming to believe that proposition; and just as one's seeing a bird fly past gives one justification for believing it did whether or not one forms this belief, adequately understanding the proposition that if the spruce is taller than the maple, the maple is shorter than the spruce, gives one (situational) justification for believing this whether one does or not.

When it comes to concepts, there is a further analogy to perception: a hierarchy analogous to the perceptual one. There is understanding a concept, such as *being taller than*. Second, there is objectually believing it to apply to something, say to a pair of things, such as the spruce and the maple. Third, there is propositionally believing something that "applies it," as when one conceives the trees as, say, the spruce and the maple, and believes that the spruce is the taller.[3]

With self-evident propositions like the straightforward proposition that if the spruce is taller than the maple then the maple is shorter than the spruce,

one need not consult one's experience of the kind of thing described, or even ponder the propositions in question, in order to grasp—roughly, to understand—those propositions. And when one does come to understand them and focuses on them in the light of that understanding, one thereby normally comes to believe and know that they are true.[4]

Two types of immediacy

There are many truths which, as just illustrated, we readily grasp and thereby immediately believe. In the simple case of comparison of heights, our belief is *immediate* in both (1) the temporal sense of 'instantly formed' and (2) the *epistemic sense*—the sense entailing that we see their truth without inferring them from anything else. The point, then, is not the temporal one that we grasp them instantly, though we may. What is crucial is that our belief exhibits *epistemic immediacy*: the belief is not based on inference or on a further, evidential belief. If it were, it would be epistemically mediate: mediated by (and thereby at least partly grounded in) the set of premises from which we infer (or on the basis of which we believe) the proposition, as my belief that Socrates is mortal is mediated by the two propositions which are part of the basis of my believing this: that he is a human being, and that all human beings are mortal.[5]

The proposition that Socrates is mortal is in another way unlike the proposition that if the spruce is taller than the maple then the maple is shorter than the spruce. It is not self-evident. There are at least two ways to explain why. First, Socrates and mortality are not intrinsically connected, as are one thing's being taller than a second and the second's being shorter than the first. An omnipotent God could have kept him in existence. Second (and speaking more generally), to know that Socrates is mortal one needs more than reflection—a temporally extended use of reason—on this proposition. One apparently needs information not given by the proposition. Even thinking of him as a human being does not absolutely preclude every route to his immortality. But reflection—indeed even *intuition* as a sometimes momentary use of reason—indicates that the spruce's being taller than the maple precludes the maple's not being shorter than the spruce.

This kind of point concerning propositions such as the one about the two trees has led philosophers to consider them to be *truths of reason*—roughly, truths knowable through the use of reason as opposed to reliance on sense experience. This kind of knowability has led philosophers to regard them as also *necessarily true—necessary*, for short, that is, such that their falsehood is absolutely precluded: there are simply no circumstances in which they are false. If a proposition is not necessary (necessarily true) *and* its negation is also not necessary, it is called *contingent*, because whether it is true—that is, its truth *or* falsity, in another terminology—is contingent on (dependent on) circumstances. That there are more than two trees in my yard is contingent. There are more, but there need not be: the number is contingent on how many I want.

The classical view of the truths of reason

How might we understand the justification of our beliefs of self-evident and apparently necessary propositions? And how do we know them? The best-known answers to these questions, and probably the only ones we should call the *classical* answers, derive largely from Immanuel Kant, though there are similar ideas in earlier philosophers who very likely influenced Kant. He discussed both the truth of the kinds of propositions in question and how we know them.[6]

What Kant said is complex and difficult to interpret precisely, and I am simply going to lay out a version of the classical account which may correspond only roughly to Kant's views. Moreover, although I am interested mainly in justification and knowledge regarding the truths of reason, I will also talk about the basis of these truths themselves when that is useful in discussing how we can know or justifiedly believe them.

Analytic propositions

Take the proposition that all vixens are female. I easily grasp its truth, and I immediately believe it: I depend on no premises or evidence. There was a time when 'vixen' was not in my vocabulary. I might then have looked at the sentence 'All vixens are female' and not known what proposition it expressed, much less seen the particular truth (true proposition) it does express. But this point does not show that I do not immediately believe that truth once I *do* (comprehendingly) consider it. It shows only that encountering a sentence which expresses a truth does not enable one to consider that truth unless one *understands* the sentence.

We can see, moreover, that when we do consider the truth that all vixens are female, we do not (or at least need not) know it on the basis of beliefs about the *sentence* 'All vixens are female'. For we can consider that same truth by using some other sentence to express it (say in Spanish), and perhaps without using a sentence at all.[7] If, however, we think about what grounds the truth of the proposition, we may discover something which in turn helps to explain why we so readily understand and believe it.

To get a sense of the ground of this truth, consider what a vixen is. It is a female fox. Indeed, the concept of a vixen may be analyzed in terms of being female and being a fox. So, in saying that a vixen is a female fox, one could be giving an elementary analysis of the concept of a vixen. Now suppose that (like Kant) we think of an analysis of a concept as indicating what the concept *contains,* or, in a certain way, *includes*. We can now say that the concept of being female is part of the concept of a vixen, and that being female is thus an element in being a vixen.[8]

In the light of all this, we might call the truth that all vixens are female an *analytic proposition*. To cite one major conception Kant presented, this is a proposition such that what it predicates of its subject can be "analyzed out of" the concept of that subject. Here the subject is vixens (or any arbitrarily

given vixen), and the predicate is *being female*, which is part of, and so analyzable out of, the concept of a vixen. The same sort of thing holds for the propositions that all bachelors are unmarried, that all triangles have three angles, that all sound arguments have true premises and true conclusions, and so on. Analytic propositions are usually considered clear cases of the self-evident.[9] This is explainable in terms of our account of the self-evident—provided we make the not implausible assumption that, given an adequate understanding of such a proposition, one can frame an analysis in which the containment relation is clearly evident. By contrast, Kant and others have viewed non-analytic propositions as *empirical*, and taken empirical propositions to be knowable not by using reason alone but only on the basis of confirmatory experience—most prominently (and perhaps necessarily) perception.

Necessary propositions

This way of looking at our example helps to explain something else that is true of the proposition that all vixens are female: it *cannot* be false and, in that sense, is necessary (a necessary truth). To see this point, try to conceive of a non-female vixen. Since the concept of a vixen is analyzable as (and hence equivalent to) that of a female fox, one is in effect trying to conceive of a *non*-female female fox. This would be both female and not female. We would have a contradiction. Hence, there cannot be such a thing, on pain of contradiction. It is thus absolutely impossible—in a sense implying impossibility by the laws of logic—that there be a non-female vixen. By contrast, it is possible that there is, and also that there is not, a 200-pound vixen. The proposition that all vixens weigh less (or more) than this is *contingent*: neither necessarily true nor necessarily false.

Because the falsity of analytic propositions entails a contradiction in this way, they are often thought to be—and are sometimes even *defined* as—those that are true *on pain of contradiction*. That is, their falsity entails a contradiction, and hence they can be false only if a contradiction is true. That is absolutely impossible. Analytic propositions are therefore regarded as truths that hold in *any* possible situation and hence are necessary (though other kinds of truths are also considered necessary).

Now if analytic propositions are true by virtue of the sort of conceptual containment relation we have been exploring, might we not know each one we do know in virtue of grasping the containment relation basic to it, in the sense that we have an adequate understanding of that relation? In considering the proposition that all vixens are female, one in some way grasps the containment relation between the concept of a vixen and that of being female. Intellectually—*intuitively*, in one widely used terminology—one sees the relation and thereby sees and (non-inferentially) knows the truth it underlies.

It might be objected that the correct account is instead this. One quickly or subconsciously reasons: The concept of a vixen is analyzable as that of a female fox; *being female* is contained in that analysis; hence all vixens are

female. So, it may be claimed, one knows that all vixens are female only inferentially. A defender of the classical view would reply that this second-order reasoning indicates how one might *show* that one knows that all vixens are female, but it does not indicate *how* one *knows* it, at least not if one just grasps its truth in the normal way.

The classical account can grant that one perhaps *could* come to know the proposition in that indirect way, by conceptual analysis. But one need not come to know it in that way; and normally, if one did not *already* know that vixens are female foxes, one would not even be in a *position* to know (on one's own) the sophisticated truth that the *concept* of a vixen is analyzable as that of a female fox. Believing that all vixens are female, in virtue of grasping the crucial containment relation between the concept of a vixen and that of a female, does not require coming to know this proposition in that sophisticated way.

The analytic, the a priori, and the synthetic

We can now see how the classical account of the truths of reason might apply to apparently *non*-analytic truths that are directly and intuitively grasped. Think about the proposition that nothing is both red and green all over at one time (different kinds of examples will be considered in Chapters 12 and 14). This is apparently self-evident and hence a truth of reason. But is it analytic? Can we analyze *being non-red* out of the concept of being green, or *being non-green* out of the concept of being red, so that anyone who said that something is red and green all over at once could be shown to be implying that it is (wholly) red and non-red, or green and non-green? This is doubtful. For one thing, it is not clear that we can analyze the concept of being red (or the concept of being green) *at all* in the relevant sense of 'analyze'. Still, on the classical view, we can know through the use of reason the necessary truth that nothing is red and green all over at once.

Let us consider two kinds of objections to the idea that the proposition that nothing is red and green all over at once is self-evident and necessary, yet not analytic. The first is based on treating the proposition as empirical and contingent; the second objection says it is analytic after all.

Take the contingency objection first. One might think that there could be a scientific explanation of why nothing is red and green all over at once; and if there is, then (on a plausible and standard view of such matters) the proposition is empirical and not self-evident or even necessary. How might such an explanation go? We can, after all, scientifically clarify what being red (or any other color) is by appeal to facts about light. This might seem to enable us to know all there is to know about basic relations among colors, even though the relevant facts about light are contingent. On the classical view, however, although scientific investigation helps us to understand certain facts about red *things* (and perhaps about the property of being red), it does not indicate what is essential to the *concept* of a red thing, such as being non-green at

the time it is red. Similarly, it is essential to the concept of a vixen that it is equivalent to that of a female fox.

To be sure, one could discover scientifically that vixens have a unique tracking system. But normally one would be identifying them for study as *female* foxes and hence would not set out to discover whether they are female. On the classical view, we cannot identify anything as a vixen—say, for experimental purposes—except under the assumption that it is female. Thus, the possibility of discovering anything inconsistent with its being female is ruled out from the start. If our experimental subject is *selected* by its having a specified property, we cannot find out experimentally that *it* (as opposed to something else it may turn into) lacks that property.

Similarly, one would not normally set out to discover scientifically whether what is red all over is ever also green all over at the same time—since it would be at best difficult to wonder whether this is true without immediately *seeing* that it is. This does not make analytic or any self-evident truths more important than scientific truths. The former are simply different: they are not of the right kind to be open to scientific verification or falsification, and in part for this reason they also do not compete with scientific truths.

It appears, then, that the suggested "scientific" objection to the classical view fails. If, however, the proposition that nothing is red and green all over at once is not a "scientific truth," that might be because it is analytic after all. Let us explore further whether the classical view is correct in claiming that the two self-evident truths in question still differ in this: being non-green is not analyzable out of the concept of being red, whereas being female is analyzable out of the concept of being a vixen.[10]

This brings us to the second objection. The objection proceeds by arguing (against the classical view) that the proposition that nothing is red and green all over at once is analytic. Could one not *indirectly* analyze the concept of being red as equivalent to the concept of having a color *other than* green and blue and yellow, and so on, in which we list all the remaining colors? This claim may seem right, because it *seems* self-evidently true that red is the only color filling that bill. But the claim is doubtful. For one thing, it is questionable whether a determinate list of all the other colors is even possible. More important, even if it is possible, the concept of being red is not *negative* in this way. To be red is to have *that color*; to be red is not simply to be a color other than green, blue, yellow, etc. Third, there is an important disanalogy to our paradigm of the analytic: whereas one could not have the concept of a vixen without having the concepts of a fox and a female, one could have the concept of being red (and so understand that concept) without even having all of these other color concepts (even if one must have *some* other color concept).

Moreover, proponents of the classical view would stress here (what is independently plausible) that an analysis does not merely provide a *conceptual equivalent*, that is, one which (necessarily) applies to the same things to which the concept being analyzed does, as the concept of being not-not-red applies

to everything the concept of being red does. An analysis of a concept (as we shall see in Chapter 11 in exploring analyses of the concept of knowledge) must meet at least two further conditions. First, it must exhibit a suitable subset of the elements that constitute the concept; second, it must do so in such a way that one's seeing that they constitute it can yield some significant degree of *understanding* of the concept. The concept of being red is surely not constituted by the complex and mainly negative property of being a color that is not green, not blue, and so on; and one could not understand what it *is* for something to be red simply in terms of understanding that long and perhaps indefinite list.

The relevant notion of understanding is *understanding of,* not *understanding that,* which is plausibly considered a special kind of knowledge of the proposition indicated by the 'that' clause, as in the case of understanding that citizenship requires being politically informed. Understanding *of* has an intimate connection with explanation. The implication of this point here is that an analysis of a concept must provide sufficient understanding of it to provide at least some *explanation* of it. The analysis of the concept of a vixen as a female fox provides material for an elementary explanation of that concept; but noting that being red is equivalent to being non-green, non-blue, and so on for all the other colors would not provide any explanation of what it is to be red. The concept of being red is simply not thus analyzable. Even the property of being red is not identical with that negative property. Indeed, one could presumably understand the list of other colors quite well even if one had never seen or imagined redness, and one had *no* perceptual, imaginational, or other concept of redness.

The point that an analysis should provide understanding of the kind that goes with explanation must not be taken to imply that we can have understanding *only* when we can explain. It is arguable, in fact, that the concept of redness is *simple* in the sense that, unlike that of a vixen, it is not analyzable into elements of any kind. One's understanding of the concept does not require its analyzability; it is enough to be able (above all) to apply it to the right things, withhold it from the wrong ones, and see what follows from its application—such as the thing's not being green.

On balance, then, it appears that the proposition that nothing is red and green all over at once is not analytic. This does not, however, prevent our rationally grasping the truth of that proposition. Truths that meet this rational graspability condition—roughly a knowability through conceptual understanding condition—have been called *a priori propositions* (propositions knowable 'from the first'), because they have been thought to be such that they can be known a priori, in a very strict sense of this phrase: known not on the basis of sense experience but simply through reason as directed toward them and toward the concepts occurring in them, at least if reason is used extensively enough and with sufficient care. Propositions that are a priori in this strict, knowability sense—as is the proposition that nothing is red and green all over at once—are also plausibly considered self-evident.[11]

Moreover, the kind of justification for believing a self-evident proposition when we believe it in the indicated way is a basic kind of justification and is often called a priori.

By contrast with analytic propositions, however, the kind of a priori proposition exemplified by that one seems to assert something beyond what analysis of the relevant concepts can show. For this reason, propositions of this kind are also called *synthetic propositions*, though these are typically defined negatively, simply as *non-analytic*. Positively conceived, they typically bring together or "synthesize" concepts and properties, even if in a negative way (as by linking redness with colors other than green—by including it among these other colors). Synthetic propositions need not, even *in part*, analyze concepts, and many are empirical in the straightforward way in which propositions evident to the five senses are.

It is noteworthy that although analytic propositions are characterized roughly in terms of how they are *true*—by virtue of conceptual containment (or, on a related account, on pain of contradiction)—a priori propositions are characterized in terms of *how they are* known, or can be known: through the operation of reason.[12] (This allows that they can *also* be known through experience, say through receiving testimony, at least if the attester's knowledge is, directly or indirectly, grounded in the operation of reason.)

On this basis, a priori propositions are also negatively characterized as knowable "independently of experience," in which this phrase above all designates no need for evidential dependence on experiential grounds, such as those of perception. But even if this negative characterization of a priori propositions is correct so far as it goes, understanding them through it will require understanding of the kinds of positive characteristics I am stressing. Let us pursue these.

Three types of a priori propositions

If we take knowability through the use of reason as a rough indication of what constitutes the a priori in general, then it includes not only self-evident propositions but certain others that are not self-evident: most clearly those propositions *not* themselves knowable simply through reason as directed toward them and toward the concepts occurring in them, but self-evidently *following from* (entailed by) such (self-evident) propositions. This is the simplest case of what is *a priori in the broad sense*. Consider the proposition that either nothing is red and green all over at once or I am flying to the moon. This self-evidently follows from the proposition about red and green, which (apparently) is self-evident. It self-evidently *follows* because it is self-evident that *if* nothing is red and green all over at once, then *either* that is true *or* I am flying to the moon.

One might think that this disjunctive (either–or) proposition is self-evident because it is so obviously both true and necessary. But even though this is true, one knows it, not in virtue of understanding the proposition

itself, but in virtue of its self-evidently following from something that *is* self-evident. One knows it inferentially, on the basis of knowing the simpler proposition that nothing is red and green all over at once. One cannot know it *just* from understanding it, as with a self-evident proposition, but only through seeing the quite different truth that if nothing is both red and green at once, then either that proposition is true or I am flying to the moon. This conditional (if–then) proposition *is* self-evident; hence, it is an utterly secure ladder on which to climb from knowledge that nothing is red and green all over at once to knowledge that either this is so or I am flying to the moon. That disjunctive proposition is *a priori in the broad sense*.

Suppose, however, that a proposition is neither self-evident nor self-evidently entailed by a self-evident proposition, but *is* provable by self-evident steps (perhaps many) from a self-evident proposition. Because there is more than one step and there can be many steps, such a provable proposition might or might not be knowable without reliance on memory, depending on the mental capacity of the rational being in question. Nonetheless, since it can be known through such a *rigorous proof*—one that begins with a self-evident proposition and proceeds only by self-evident steps (entailments) to its conclusion—a rigorously provable proposition may be called *ultimately a priori* (or ultimately self-evident, though the former term seems preferable). It is not a priori in the broad sense because (1) it is not linked to the self-evident by a *single* step and—more important—(2) it is not necessarily self-evidently linked to it.[13] But as it is ultimately traceable to a self-evident proposition, it may be considered a priori in the ultimate provability sense.

Thus, in speaking of propositions that are a priori in the most comprehensive terminology, I include not only the intuitively central cases that are self-evident or just one step from it—propositions self-evidently entailed by a self-evident proposition—but also those not thus entailed but nonetheless provable by self-evident steps from a self-evident proposition.

We could say, then, that for the kind of classical view in question, the self-evident is the *base* of the a priori: a priori propositions are those that are either self-evident (i.e., a priori in the narrow sense) or, though not themselves self-evident, self-evidently follow from at least one proposition that is (hence are a priori in the broad sense). The general notion of an a priori proposition, applicable to both cases and others, is roughly the notion of a truth that either is a self-evident proposition or is self-evidently entailed by one, or provable from one by self-evident steps.[14]

Knowledge of propositions a priori in the broad or ultimate provability sense, unlike knowledge of those a priori in the narrow sense, depends on knowledge of some self-evident proposition as a ground. But neither kind of knowledge depends on knowledge of any empirical proposition, and in that sense both kinds are "independent of experience."

It is because a priori propositions (of any sort) are understood in relation to how they can be *known* that the notion of the a priori is commonly considered epistemological. But many a priori propositions also have a special

property of a different kind. Many are said to be analytic. The notion of the analytic is more often taken to be of a different, non-epistemological kind, say conceptual, since analytic truths are conceived as grounded in a simple containment relation between concepts.[15]

It should not be surprising, then, that the categories of the analytic and the a priori are not identical. In both cases, however, proponents of the classical view have taken the relevant propositions to be necessary: this is commonly thought to be obvious for the analytic ones, which are true "on pain of contradiction," but it has seemed reasonable to classical theorists to hold that even synthetic a priori propositions are invariably necessary. Perhaps the underlying thought is that if their truth were contingent and so depended on what holds in (is contingent on) some possible situations but not others, one could not know it just on the basis of understanding the proposition itself.

The empirical

A huge variety of truths are not a priori. That the spruce is taller than the maple is one of them. Truths that are not a priori are called *empirical* (or *a posteriori*) *truths*. This means, roughly, that the propositions in question can be known only *empirically*: knowable (assuming they are knowable) only on the basis of experience, as opposed to reason—above all on the basis of perceptual or self-conscious experience (in the ways described in Chapters 1, 2, and 4).

Saying simply that a proposition is empirical (or a posteriori) leaves open whether it is true: there are empirical falsehoods, such as that it is not the case that the spruce is taller than the maple, as well as empirical truths. (In this the term 'empirical proposition' is unlike 'a priori proposition' and 'necessary proposition', which are not commonly used to refer to falsehoods, but my main examples of empirical propositions will be truths.)

For the classical view, empirical propositions as well as a priori propositions are crucial for our lives. Indeed, the former include every truth known perceptually, such as those known through observing the colors and shapes of things, and all truths known scientifically, such as generalizations linking the temperatures and the volumes of gases, or ingestion of drugs with change in behavior. A certain range of a priori propositions, such as those of logic and pure mathematics, are presupposed by both common sense and science. Empirical propositions are also required to guide us in dealing with the world, but the classical view sees them as open to disconfirmation through experience in a way that a priori propositions are not.

Analytic truth, concept acquisition, and necessity

Analytic truths, as well as certain synthetic ones, are called a priori because analytic truths are knowable through the use of reason. But analytic truths appear to be knowable—or at least are showable—through a different use

of reason than is appropriate to the synthetic a priori truths. It may be that I know that nothing is red and green all over at once by virtue of simply grasping, as a rational being, a kind of incompatibility between the concept of being red (at a time and place) and the concept of being green. But, as pointed out earlier, I apparently do not know it by virtue of grasping a *containment* relation between being red (or green) and anything else. If this does not illustrate two different uses of reason, it at least indicates a different kind of application of reason to different kinds of relations of concepts.

Because my knowledge of the proposition that nothing is red and green all over at once is not based on grasping a containment relation, it differs from my knowledge of the analytic truth that all vixens are female. Yet in both cases the relation between the concepts involved in the truth seems to be the basis of that truth. In both, moreover, I apparently know the truth through rationally *understanding* that relation: a relation of analytic containment in one case, and of mutual exclusion in the other.

These points do not imply that experience is irrelevant to knowledge of the a priori. On the classical view, I do need experience to *acquire* the concepts in question, for instance to acquire color concepts or the concept of a fox. But once I have the needed concepts, it is my grasp of their relations, and not whatever experience I needed to acquire the concepts, which is the *basis* of my knowledge of analytic and other a priori truths.

In part because of these similarities, as well as because the falsity of a priori propositions seems inconceivable, the classical view takes synthetic a priori truths as well as analytic truths to be necessary. They cannot be false, even though in the synthetic a priori cases it seems not to be strictly contradictory to deny one. For instance, claiming that something *is* red and green all over is not contradictory in the sense that it (formally) entails that some proposition—say, that the object in question has a definite color—is and is not true. Still, on the classical view it is absolutely impossible that something *be* red and green all over at once. We need only reflect on the relevant concepts (mainly the color concepts) to realize that nothing is red and green all over at once; we readily grasp (apprehend) an exclusion relation between being red and being green.

It is also commonly held by philosophers in the classical tradition that all necessary propositions are a priori. One rationale for this might be that necessity is grounded in relations of concepts and these (or at least the relevant relations) are the same in all possible situations. A mind that could adequately survey all possible situations (like the divine mind as often conceived) could thus know the truth of all necessarily true propositions. Since this survey method would be possible without analyzing one concept out of another, the grounding of necessity in conceptual relations would also explain how there can be synthetic necessary truths. And for the classical view, these, being necessary, are also a priori.[16]

Summarizing, then, the classical view says that all necessary propositions are a priori and vice versa, but it maintains that analytic propositions

constitute a subclass of a priori ones, since some a priori propositions are synthetic rather than analytic. The view tends to conceive the truth of all a priori propositions as grounded in relations of concepts (or of similar abstract entities, such as "universals," in Bertrand Russell's terminology).[17] But the position conceptually accounts for these propositions differently: for necessary propositions in terms of the unrestricted *circumstances* of their truth (the absolute impossibility of their falsehood in any circumstances), for analytic ones in terms of *how* they are true (typically, by virtue of containment relations), and for a priori propositions in terms of how their truth is *known* (through understanding).

The empiricist view of the truths of reason

The classical view of the nature of what I am calling *a priori truths—truths of reason*—and of our knowledge of them has been vigorously challenged. To appreciate the epistemological significance of reason as a source of justification and knowledge, and of truths of reason themselves, we must consider some alternative accounts of these truths.

John Stuart Mill held that ultimately there are only empirical truths and that our knowledge of them is based on experience, for instance on perception.[18] We might call this sort of view *empiricism about the* (*apparent*) *truths of reason*. The name suits the view, since the position construes apparently a priori truths as empirical, though it need not deny that reason as a capacity distinct from perception has *some* role in giving us justification and knowledge. Reason may, for example, be crucial in extending knowledge by enabling us to prove geometrical theorems from axioms. But the view I want to explore (without following Mill in particular) denies that reason grounds justification or knowledge in the non-empirical way described by the classical theory.

Rationalism and empiricism

Before we consider Mill's thesis in detail, we should contrast it, from the most general epistemological point of view, with that of Kant and other rationalists to get a better sense of what is at stake in the controversy between rationalism and empiricism. Kant's position on the truths of reason might be called rationalist, Mill's empiricist. These terms are used too variously to make precise definition wise. Very roughly, however, *rationalism* in epistemology takes reason to be far more important in grounding our knowledge than empiricism allows, and rationalists virtually always assert or imply that, in addition to knowledge of analytic truths, there is knowledge of synthetic a priori truths. Very roughly, *empiricism* in epistemology takes experience, most notably sensory experience, to be the basis of all of our knowledge except possibly that of analytic propositions, understood as including purely logical truths, such as the truth that if all whales are mammals and no fish

are mammals then no whales are fish. (For both empiricists and rationalists, analytic propositions are typically taken to include logical truths.)[19]

One might wonder why some empiricists grant that analytic truths may be a priori. The central point (though an empiricist might not put it this way) may be seen if we use the terminology of the classical theory: even if such logical propositions are not true by virtue of containment relations between concepts, their negations formally entail contradictions, for instance that some vixens are and are not female foxes. They are therefore paradigms of truths of reason; for the use of logic alone, which is perhaps the purest use of reason, can show that they can be false only if a contradiction is true—which is absolutely impossible. This is another reason why, as noted above, analytic propositions are sometimes given a broader characterization than I have proposed and are taken to be those whose negations entail a contradiction.[20]

Some empiricists do not allow that any knowledge, even of so-called analytic propositions, is genuinely a priori. A *radical empiricist*, such as Mill, takes *all* knowledge to be grounded in experience. A *radical rationalist* (which Kant was not) would take all knowledge to be grounded in reason, for instance to be intuitively grounded in a grasp of self-evident propositions or deductively based on inference from a priori truths that are intuited.[21]

Empiricism and the genesis and confirmation of arithmetic beliefs

Empiricism about what are called the truths of reason is most plausible for the apparently synthetic a priori ones, so let us sketch it with reference to an apparently synthetic kind of a priori proposition that has been much in dispute. Mathematical truths, particularly truths of simple arithmetic, are often regarded as synthetic a priori. Consider the proposition that $7 + 5 = 12$ (Kant's example, also found in Plato's *Theaetetus*). It is easy to say that one just knows this, as one knows that nothing is red and green all over at once. But how does one know it?

Here we cannot readily find a good analogy for the simple exclusion relation we apparently grasp in the case of red and green. Could it be that from experience with objects, say with counting apples, then combining two sets of them, and recounting, we learn our first arithmetic truths and then use reason to formulate general rules, such as those for calculating larger sums?

Viewed in this way, arithmetic develops rather as scientific hypotheses are often thought to, with observations crucial at the base, generalizations formulated to account for them, and broader generalizations postulated to link all the observations and the narrower generalizations together. And do we not first learn to add by counting physical things, or by counting on our fingers?[22]

To be sure, we perhaps cannot imagine how the number 7 added to the number 5 could fail to equal the number 12. But that is not a point about the behavior of objects in the physical world. The physical world *could* go

haywire so that when (for instance) five apples and seven oranges are physically combined, the result of counting the new set is *always* eleven. If that happened pervasively, might we not begin to think that arithmetic must be revised, just as Einstein's work showed that the physics of "the incomparable Sir Isaac Newton" needed revision? Perhaps the crucial epistemological consideration is what overall account of our experience is most reasonable; and if the best overall account should require rejecting a proposition now considered a priori and necessary, so be it.

The classical view provides for several critical responses. One concerns the distinction between two quite different things: the *genesis* of one's beliefs—what produces them—and their *justification*, in the sense of what justifies them. A second point concerns whether arithmetical propositions can be tested observationally. A third focuses on the possibility of taking account of what looks like evidence against arithmetical truths, so that even if one's final epistemological standard for judging a proposition *is* its serving the demands of the best overall account of experience, these truths can be preserved in *any* adequate account. Consider these ideas in turn.

First, granting for the sake of argument that our arithmetic beliefs arise from counting physical objects, is the experience that produces them what *justifies* them? The genesis of a belief—what produces it—is often different from what justifies it. The testimony of someone I realize is unreliable might, when I am off guard, produce my belief that different brands of aspirin do not, apart from additives, differ chemically. My belief would at that point be unjustified; but it might become justified later when I learn that aspirin is simply acetylsalicylic acid. Moreover, regardless of what produces our arithmetic beliefs initially, when they are justified in the *way* my belief that $7 + 5 = 12$ now is, experience does not appear to be what justifies them. For my part, I do not see precisely how the truth of the proposition might be grounded in the behavior of objects when they are combined; and I would not try to justify it, as opposed to *illustrating* it, by citing such behavior.

This brings us to the second point: it is doubtful that the proposition that $7 + 5 = 12$ is (empirically) *testable*, say by examining how objects combine, though it is *exemplifiable* in that way. The empiricist might reply that this by no means shows that the proposition is, as the classical view insists, necessarily true rather than contingent and empirical. Indeed, it does not. But let us look closely at the idea that it could be tested, and could thereby be disconfirmed by discovering that when groups of five objects are combined with groups of seven, we find just eleven.

This brings us to a third response. How might one deal with repeated and systematic counter-evidence? Classical theorists will argue that it is possible for the world to alter in such a way that this combination procedure results in one item's disappearing, or in our failing to see it, or in our misremembering how many items entered the mix before our re-counting. They will also argue that the unexpected realization of such possibilities would be a better interpretation of the strange cases described—hence of our overall

experience—than saying that it has turned out to be false that 7 + 5 = 12. Thus, instead of saying that an arithmetical principle has been falsified, we would say that the world no longer uniformly exemplifies it.

One consideration favoring the classical view is that it is at best difficult even to understand how the purely arithmetical principle could be false. The number 7 plus the number 5 apparently equals the number 12, regardless of how apples and oranges behave. For the arithmetic statement is apparently not *about* apples and oranges, though (so far as we know) their behavior exemplifies it. For the classical view, at least, it is about *numbers*, which, unlike the arabic or roman or other *numerals* we use to represent them linguistically, are abstract and non-physical. If a proposition is not about concrete objects, facts about their behavior are not a test of its truth.

Notice something else. In order to gather purportedly significant counter-evidence to the arithmetic proposition in question, one would have to rely, as already noted, not only on memory and perception (both highly fallible sources) but also on simple arithmetic: one would have to *count* disconfirming cases. A single apparent instance, say of seven and five things brought together and not adding up to twelve, would not be significant, and one must keep track of how many anomalies there are, relative to confirmatory instances in which the expected sum is counted out. It is not normally reasonable to give up a good theory on discovering a single apparent counter-instance. It appears, then, that in order to take seriously empirical evidence that would undermine arithmetic, we must trust perception in our counting, arithmetic itself in our summing, and memory in our overall judgment.

One might think it is enough simply to *have* a significant number of such disconfirming cases. But this is not so. One must be justified in believing that the number *is* significant. And how could one achieve this if one either made no count or—in any case—could not rely on one's count of single cases to sum to a significantly large number? If it need not be true that 7 + 5 = 12, why should 1 + 1 + 1 disconfirming instances necessarily sum to 3? And would anything less than a huge number of apparently disconfirming cases be evidentially decisive against such a proposition of simple arithmetic? A single disconfirming instance would surely seem just an anomaly; there must be a significant number. One would, then, have to rely on some arithmetic propositions, such as that 1 + 1 + 1 disconfirmations = 3 (a minimally significant number, perhaps), in order to mount an effective challenge to the (necessary) truth that 7 + 5 = 12. Given the interconnections among arithmetic propositions, it is not clear that one could consistently (or at least with any plausibility) maintain the needed disconfirmatory propositions while denying that 7 + 5 = 12. Still another obstacle to recognizing apparent counter-evidence as genuine is the dependence on memory to keep track of disconfirming instances. The fallibility of memory would defeat confidence that one had adequately tracked apparent disconfirmations.

There may be a way around these objections, but even finding it would leave one far from a strong case for the contingent or empirical status of

arithmetic truths.[23] Even if one appealed, not to apparent counter-instances to the proposition that 7 + 5 = 12, but to a well-confirmed *theory* to argue that it could be false, one would need to do at least some counting of one's confirmatory data regarding that theory (not to mention other ways in which theory confirmation relies on arithmetic, perception, and memory).

None of these points requires us to deny that there *is* a similar, contingent arithmetic proposition about apples and oranges, namely that when we *count* five of the first and *place* them next to the result of counting seven of the second, we *can count* twelve all told. This proposition may easily be confused with its pure mathematical counterpart. The former is clearly contingent and empirical, but its being so does not show that the purely arithmetic proposition is also. The distinction between pure and applied mathematics can also be brought to bear on geometry.[24]

There is a related metaphysical dimension of the question of the status of arithmetic truths. By contrast with at least one form of the classical view, radical empiricism denies that there *are* abstract entities and so, believing that mathematical propositions are about something concrete, radical empiricists naturally view them as generalizations about the behavior of physical objects. We need not accept the empiricist view to grant that if physical things did not exemplify the proposition that 7 + 5 = 12, the proposition would be of far less *value* to us even if necessarily true. If the physical world went haywire, it could turn out to be false that when seven apples are placed together with five more and the total collection is counted, the count yields twelve. This chaotic situation would falsify the *physical* principle already contrasted with the arithmetic one in question. But the physical principle is not, and does not even follow from, the purely mathematical proposition we are discussing.

Empiricism and logical and analytic truths

The empiricist view of the a priori can also be applied to analytic propositions and even to self-evident logical truths, and it may indeed appear more plausible in that case. Suppose that through scientific investigation we discover that vixens have certain characteristics we think of as male, such as certain hormones. Imagine that gradually (perhaps because of chemicals in the environment) these discoveries mount up so that the female foxes in our laboratory begin to seem more aptly classified as male than as female. Could not a time come when we begin to doubt that vixens are female after all?

And what about the logical principle of the excluded middle, which says that every proposition is either true or false? Consider the proposition that Tom is bald. Must this proposition be either true or false no matter what the quantity or distribution of hair on his head? Surely the proposition is an appropriate counter-example to the principle of the excluded middle.[25]

The classical view can offer its own account of these examples. For one thing, particularly over a long time, we can begin to use a term in a sense different from the one it now has. Thus, the discoveries about vixens could result in our someday using 'vixen' to mean not 'female fox', but 'fox with

female external sexual characteristics and of the anatomical kind *K*' (where *K* is the kind of animal we have in our laboratory). Then, when we utter such words as 'Vixens are not really female', we are not denying the analytic proposition *now* expressed by 'All vixens are female'. We have confirmed something else, rather than disconfirming this.

In this way, then, our experience might result in our someday no longer assertively uttering 'Vixens are female' to say anything that we believe. This certainly does not show that experience might falsify the proposition we *now* affirm when we assertively utter that. Given what we now mean by 'vixen', in saying that all vixens are female we do not rule out that *those* 'vixens' in the lab could have internal biological and chemical characteristics in the light of which *they* ultimately need not be considered female.

Regarding the principle of the excluded middle, I would stress that Aristotle plausibly argued against it, and some contemporary philosophers of logic do, too. The main reasons for doubting it, moreover, do not depend on empiricism. Let us explore some of them.

Consider again the vague statement that Tom is bald. It may certainly be argued that this need not be either true or false. It is not as if 'bald' meant, say, 'having fewer than 500 hairs on the top of one's head'. It does not. And if it did, the term 'top' would still be vague and would cause the same trouble: it would be unclear in what *area* we must find 500 hairs. If the middle possibility—neither truth nor falsity—is to be ruled out, it must be by a better argument. The principle of the excluded middle, though often used to suggest that even logical truths are not necessarily true, is controversial among rationalists and empiricists alike. The principle is a poor example to support the empiricist case against the necessity of logical truths.

When, by contrast, standard examples of simple logical truths are used, the effect seems very different. Consider the proposition that if Ann is coming by bus or she is coming by plane, and it is false that she is coming by bus, then she is coming by plane (which exemplifies the general logical truth that if at least one of two propositions is the case and the first is false, then the second is true). Is there any plausibility in the view that this might be false? I find none; and while nothing said here proves that the empiricist account of the a priori is mistaken, it appears less plausible than the classical account.

If what we have seen so far is accepted, the classical view of the truths of reason is quite defensible and the empiricist critique of it fails. But we have not yet adequately taken into account the ways in which knowledge of those truths might depend on language. This is an important topic particularly given the extent to which understanding the a priori is connected with understanding language. The next chapter will consider this topic in some detail.

Notes

1 *Adequacy* of understanding of a proposition cannot be merely partial understanding, and it is more than simply getting the general sense of a sentence expressing it, as where one can analyze the grammar of the

sentence, indicate something of what it means through examples, and perhaps translate it into another language one knows well. Adequacy here implies not only seeing what the proposition says but also being able to apply it to (and withhold its application from) an appropriately wide range of cases. This matter is treated in some detail in my 'Self-Evidence', *Philosophical Perspectives* 13 (1999), 205–28. Note also that there is no appeal here to understanding or positing the *necessity* of the propositions (though the characterization lends itself to taking them to be necessary). In this respect my notion of the self-evident is simpler and more moderate than the traditional one common in much of the literature. See, for example, Laurence BonJour, 'Toward a Moderate Rationalism', *Philosophical Topics* 23, 1 (1995), 47–78, esp. section 3.

2 For a helpful discussion of obviousness related to (but quite different from) the one in my 'Self-Evidence' and connected with the theory of the a priori in general, see Robin Jeshion, 'On the Obvious', *Philosophy and Phenomenological Research* 60, 2 (2000), 333–55.

3 Two points are appropriate here. (1) A fourth case is one in which a concept is not only *exercised* in a belief but explicitly figures in it, as when one believes that the concept *being taller than* is instantiated by the spruce and the maple. (2) The analogy between perception and conception I am developing is meant to leave open what concepts are and what it is to understand one. As will later be apparent, philosophers differ in their understanding of the truths of reason in part because of their different understandings of the nature of concepts.

4 One reason for the normality qualification is to make room for the possibility that one can consider and adequately understand a self-evident proposition yet fail to believe it. Brain manipulation might cause such failure. We should also make room for the possibility that, especially with more complex self-evident propositions—say that if *p* entails *q* and *q* entails *r* and *r* entails *s*, and *s* is not true, then *p* is false—it may take a person time to form the belief.

5 *Temporal immediacy*, unlike epistemic immediacy, is a property not primarily of beliefs as such but of their formation. A belief is temporally immediate when its formation occurs "without delay" upon the person's considering the proposition in question (or encountering the situation, such as the sight of a lightening bolt, that gives rise to the belief). One could also say that propositions are temporally immediate in a derivative sense when they are so obvious that one normally believes them immediately on (comprehendingly) considering them. Many self-evident propositions are like this. But when I consider some self-evident propositions, such as that if there never have been siblings, then there never have been first cousins, it may or may not take me a moment to see their truth. Still, when one does see such a truth, the belief one forms will (at least normally) be epistemically immediate, not inferential. So, this proposition and my coming to believe it may or may not be temporally

immediate. By contrast, the proposition that I am now seeing print is temporally immediate (for me) but is not self-evident. It is evident not in itself, but *through* what I see.

6 Kant's most detailed presentation of his views on these matters is in his *Critique of Pure Reason* (first published in 1781), but a short presentation is provided in the Preamble to his *Prolegomena to Any Future Metaphysics* (1783). Kant's conception of the analytic is quite reminiscent of Aquinas's idea that the self-evident has "its predicate contained in the notion of the subject" (*Summa Theologiae,* Question 94, Article 2).

7 There has long been controversy about whether such thought is possible without using language, or at least having a language. Donald Davidson is among those to argue for a strong dependence of thought on language. See, for example, his *Inquiries into Truth and Interpretation* (Oxford: Oxford University Press, 1984). Relevant critical discussion of Davidson is provided by Ruth Barcan Marcus in 'Some Revisionary Puzzles about Belief and Believing', *Philosophy and Phenomenological Research*, supplement to vol. 50 (1990), 133–53, which brings out serious problems for the view that beliefs must have sentence-like objects. There is no need to take a stand on this issue for my main purposes in this book.

8 *One* way to conceive this is as follows: if the concept of *F* is part of the concept of *G*, then having the property (of) *F* is self-evidently entailed by having the property (of) *G*. I do not accept this overall conception of conceptual containment but do believe that the entailment holds (even if not self-evidently).

9 This is plausible if (1) the correct analysis of a key concept in an analytic proposition, say that of a vixen, is discoverable, without reliance on anything beyond understanding that concept, by anyone with an (adequate) understanding of the proposition, and (2) given a correct analysis of that concept, the truth of the analytic proposition is appropriately evident. However, some analytic propositions are not understandable in this way; some might be *provable* only by a lengthy process from one that is (a notion discussed on page 113). Further, it is by no means clear that every analytic proposition is self-evident in the very common sense that implies a fairly high degree of obviousness. If, as seems plausible, the self-evidence of a proposition simply implies that *some* kind of adequate understanding is sufficient for justification for believing it, then we might plausibly distinguish between the immediately and the mediately self-evident and allow that the latter propositions may be understandable (to normal persons) only on the basis of considerable reflection. Cf. Thomas Aquinas's view (which Kant might have known) that:

> Any proposition is said to be self-evident in itself, if its predicate is contained in the notion of its subject . . . *Man is a rational being*, is, in its very nature, self-evident, since he who says *man* says *a rational being;* and yet to one who does not know what a man

> is, this proposition is not self-evident . . . some propositions are self-evident only to the wise, who understand the meaning of the terms of the propositions.
>
> (*Summa Theologiae,* Question 94, Article 2)

This seems to anticipate Kant's containment notion of the analytic and largely accords with the conception of the self-evident I have introduced.

10 There are philosophers who regard colors as subjective in a way that might seem to undermine the example here. I do not see that taking the proposition that nothing is red and green all over at once to be necessary, synthetic, and a priori entails any particular analysis of color properties, and I doubt that the example fails. If the example should depend on a mistaken realist account of color and for that reason fail, anti-realism about shape properties is less plausible, and the proposition that nothing is round and square might serve as well. For accounts of the status of color see C.L. Hardin, *Color for Philosophers, Unweaving the Rainbow* (Indianapolis: Hackett, 1988), and Edward Wilson Averill, 'The Relational Nature of Color', *Philosophical Review* 101 (1992), 551–88. For a detailed discussion of color properties, with application to the apparently synthetic a priori proposition that nothing is red and green all over at once and with a defense of the view that color properties supervene on (and so are determined by) dispositional properties of physical objects, see Colin McGinn, 'Another Look at Color', *Journal of Philosophy* 93, 2 (1996), 537–53.

11 This allows that such propositions can *also* be known empirically, say through testimony, though there are restrictions (discussed in Chapter 7) on how this may occur. The characterization suggests that an a priori proposition is knowable non-inferentially even if only on the basis of considerable reflection, but the exact mode of the appropriate reflection is not something that need be settled here. A full account of this conception of the a priori would explicate the kind of possibility of knowledge in question; it is presumably not mere logical possibility in the sense that no contradiction is formally entailed by the occurrence of the relevant knowledge, but a conceptual possibility, roughly in the sense that such knowledge is provided for by the concept of the relevant kind of knowledge: the kind grounded in understanding propositions of the sort in question. My preference is to characterize the a priori in terms of self-evident propositions and leave open what kind of possibility there has to be of the sort of understanding that grounds justification for believing those propositions. For a valuable treatment of possibility and necessity arguing that such modal notions are irreducible, see Scott A. Shalkowski, 'Conventions, Cognitivism and Necessity', *American Philosophical Quarterly* 33 (1996), 375–92.

12 Kant's section 2b of his Preamble to the *Prolegomena to any Future*

Metaphysics (trans. by Lewis White Beck, New York: Liberal Arts Press, 1950) opens with 'The Common Principle of All Analytical Judgments is the Law of [non]Contradiction' and almost immediately continues: "For the predicate of an affirmative analytical judgment is already contained in the concept of the subject, of which it cannot be denied without contradiction."

13 There is a subtlety here that needs comment. Imagine that a self-evident axiom, *A*, self-evidently entails a theorem, *t*, which in turn self-evidently entails a second theorem, *t′*. Self-evident entailment (as opposed to entailment in general) is not *transitive*: *A* can self-evidently entail *t* and *t* can self-evidently entail *t′* without *A*'s self-evidently entailing *t′*. Here one could understand the conditional proposition that if *A* then *t′* quite adequately without thereby having justification for believing it. One might need the intermediate step, *t*, to achieve that justification, and it need not be discerned simply in adequately understanding the conditional itself. This possible limitation does not preclude there being *some* kind of understanding of that conditional and related concepts, such as a perfectly omniscient being might have, in virtue of which the proposition that if *A* then *t′* can be seen to be true. This shows that—as Aquinas saw in the quotation from him in note 9—there is a related notion—self-evidence *for a particular person* (or mind)—which must be distinguished from self-evidence in its basic, non-relativized form, making reference only to anyone's understanding. Still, even if what is self-evident for God might not be self-evident for us, some propositions are unqualifiedly self-evident. The case also shows that not every proposition *provable* by individually self-evident steps from a self-evident premise may be assumed to be a priori in the (moderately) broad sense of being self-evidently entailed by a self-evident proposition; for (as just explained) such a proposition might not be self-evidently entailed by a self-evident proposition.

14 In a broader usage, a falsehood can be called an a priori proposition provided it is an a priori *truth* that it is false. This less common usage raises no special problems but presents a terminological complication I ignore in the text.

15 There is much difference in judgment about how to classify the analytic. It might be considered a semantic concept by those who think of it as truth by virtue of the *meanings* of the relevant terms. It might be regarded as ontological by those who think such truths are basic to the structure of reality. For epistemology, the notion of the a priori is the more important of the two. For an immensely influential paper arguing that neither notion is clear see W.V. Quine, 'Two Dogmas of Empiricism', in his *From a Logical Point of View* (Cambridge, MA: Harvard University Press, 1953). Among the widely noted replies is H.P. Grice and P.F. Strawson, 'In Defense of a Dogma', *Philosophical Review*

55 (1956), 114–58. For more recent discussion of these issues see Gillian Russell, *Truth by Virtue of Meaning: A Defense of the Analytic/Synthetic Distinction* (Oxford: Oxford University Press, 2008).

16 Critical discussion of the question whether the a priori must be necessarily true is provided in my 'Skepticism about the A Priori: Self-Evidence, Defeasibility, and *Cogito* Propositions', in John Greco (ed.), *The Oxford Handbook of Skepticism* (Oxford: Oxford University Press, 2008).

17 See Bertrand Russell, *The Problems of Philosophy* (Oxford: Oxford University Press, 1912), Chapters 8–10 (these chapters are reprinted in Huemer, *Epistemology*).

18 See especially J.S. Mill, *A System of Logic* (first published in 1843), particularly Book II, Chapters 5–7. For a much more sophisticated critique of apriorism in mathematics and an empiricist account of mathematical truths, see Philip Kitcher, *Mathematical Knowledge* (Oxford: Oxford University Press, 1984).

19 Granting it is at best not obvious how logical truths are knowable by any analysis that reveals containment relations, their negations can be clearly seen to entail contradictions.

20 How broad this is depends on the notion of entailment used. I have in mind a notion for which the negation of a proposition entails a contradiction provided that the use of formal logic, supplemented only by (correct) definitions, renders a contradiction deducible.

21 Someone might think all truth is a priori on the ground that it is true a priori that (1) God exists; (2) a certain universe specifiable in every detail is the best of all possible universes; and (3) God creates the best of these universes. Then, with sufficient intellectual power, one could (arguably) reason one's way to any truth. Gottfried Wilhelm Leibniz (1646–1716) has been read as holding a view close to this (but there are reasons to doubt that he did, including considerations about divine freedom).

22 Cf. W.D. Ross, explicating how Aristotelian *intuitive induction* can yield a priori knowledge: "We find by experience that this couple of matches and that couple make four matches . . . and by reflection on these and similar discoveries we come to see that *it is of the nature of two and two to make four*" (*The Right and the Good* [Oxford: Oxford University Press, 1930], p. 32).

23 The proposition that $1 + 1 + 1 = 3$ might be held to be more intuitive than the proposition that $7 + 5 = 12$. But, first, in practice we might need to rely on less intuitive or much more complicated arithmetic to get a good case for the possible falsehood of the original proposition; second, and more important, the simpler proposition that $1 + 1 + 1 = 3$ will also do as a case of a necessary mathematical truth.

24 For discussion of the status of the a priori in connection with geometry, see the Appendix to Laurence BonJour, *In Defense of Pure Reason* (Cambridge: Cambridge University Press, 1998). That book is also of

interest for its criticism of Kant, who in BonJour's view is less a rationalist about—and less plausible concerning—the a priori than is often thought.

25 For discussion of vagueness and its bearing on epistemological matters (as well as references to his own and others' earlier work on vagueness) see Timothy Williamson, *Vagueness* (London: Routledge, 1994) and *Knowledge and its Limits* (Oxford: Oxford University Press, 2000).

6 Reason II

Meaning, necessity, and provability

- **The conventionalist view of the truths of reason**
  Truth by definition and truth by virtue of meaning
  Knowledge through definitions versus truth by definition
  Conventions as grounds for interpretation

- **Some difficulties and strengths of the classical view**
  Vagueness
  Meaning change and falsification
  The possibility of empirical necessary truth
  Essential and necessary truths
  Necessity, apriority, and provability

- **Reason, experience, and a priori justification**
  A priori beliefs
  Loose and strict senses of 'a priori justification' and 'a priori knowledge'
  The power of reason and the possibility of indefeasible justification

6 Reason II

Meaning, necessity, and provability

The radical empiricist critique of rationalism is neither the only kind empiricists can mount nor the only plausible source of objections to it. Another important approach to understanding the truths of reason and our justification and knowledge builds on the undeniable connections between how we use our language—specifically, on our linguistic conventions—and our knowledge of truths expressible in that language.

The conventionalist view of the truths of reason

We have seen the importance of analyses for understanding the a priori. Definitions of certain kinds may be considered linguistic counterparts of analyses. On one view, analytic truths may be better seen as definitional than as "analytical." This idea needs examination.

Truth by definition and truth by virtue of meaning

To see how this approach goes, suppose that analytic propositions may be said to be *true by definition*. On the assumption that the truth or falsity of definitions turns on linguistic conventions, one can now make moves parallel to the classical ones that are expressed in terms of concepts. Thus, 'vixen' is definable as meaning (the same thing as) 'female fox'; 'female' is part of this phrase; hence, by grasping a definition (even if we do not call it to mind) we can *see* how the proposition that all vixens are female is true. The predicate, 'is female', expresses part of the meaning of the subject, 'vixen', just as the concept of being female is part of the content of the concept of a vixen. Thus, according to conventionalism, by *appeal* to the definition of 'vixen' as having the same meaning as 'female fox', we can also *show* that the proposition that all vixens are female expresses an analytic truth.

The conventionalist may grant that in the case of synthetic truths of reason, for instance the truth that nothing is red and green all over at once, we cannot make the same moves. For the relevant color terms are indefinable, or in any case not definable in the needed way. But we can still speak of truth by virtue of *meaning* or at least *convention*, in the limited sense that it seems

to be a matter of the meanings of, or conventions governing, say, the terms 'red' and 'green', that if one of the terms applies to a surface at a time and place, the other does not. Why else would someone who sincerely denies that nothing is red and green all over at once seem to exhibit an inadequate understanding of at least one crucial term used in expressing that proposition?

What terms mean is a matter of convention. It depends entirely on agreement, usually tacit agreement, among the users of the relevant language, concerning the proper application of the term. We could have used 'vixen' differently; we in fact would have done so if the history of our language happened to differ in a certain way. Moreover, even now we could decide to use 'vixen' differently and proceed to do so.

The suggested account of the truths of reason—*conventionalism*—grounds them in conventions, especially definitional conventions, regarding meaning. Secondly, and related to this basic claim, it conceives our knowledge of them as based on our knowing those conventions. Since knowledge of conventions is reasonably taken to be empirical knowledge based on suitable observations of linguistic behavior, conventionalism (on this interpretation) turns out to be a kind of empiricism regarding the truths of reason, and it has been held by some philosophers in the empiricist tradition. The claim is not that these truths are *about* words, but that knowledge of them is *based* on empirical knowledge of linguistic usage.

Knowledge through definitions versus truth by definition

Some of the points made by conventionalism are quite plausible. In grasping the definition of 'vixen' as meaning the same thing as 'female fox', perhaps we can see that all vixens are female; and under certain conditions, by appeal to the definition we perhaps can show that this truth holds. But do these points undercut the classical view? If the points hold, that may well be because of something *non*-linguistic: perhaps, *in* grasping the definition we understand the *concepts* involved and thereby see a containment relation between the concept of a vixen and that of being female. In this or some other way, understanding definitions might be a ladder by which we climb to an understanding of concepts.

Furthermore, as a proponent of the classical account might also note, it seems possible to grasp the relevant conceptual relations, and thereby already know the analytic truth, even if one does not know any such definition. Indeed, it might be only on the basis of the analytic truths one knows—such as that all vixens are female, and that all female foxes are vixens—that one is able to *construct* a definition of 'vixen'—with its present meaning—in the first place. The definition would reflect what is already true in virtue of how the concepts in question are related; the concepts are not themselves created by or grounded in linguistic conventions.

Contrary to conventionalism, then, the knowledge of analytic truths would be essential in one's route to the definitional knowledge, not the other

way around. Understanding the relations between the concepts expressed by the words in question would be the basis for judging the definitions of those words; it would not be through first knowing the truth of those definitions that one understands the conceptual relations or knows the analytic truth. Hence, knowledge of analytic truths apparently does not *depend* on knowledge of definitions or conventions, even if the former can sometimes be gained through the latter.

The more general important point implicit here is that conventionalism fails to give a good account of what grounds the *truth*, as distinct from our knowledge—or some of our knowledge—of analytic propositions. It is not *because* 'vixen' means the same thing as 'female fox' that all vixens are female. For, as we saw in assessing the empiricist view, this analytic truth does not depend on what 'vixen' means. This truth holds whether there is such a word or not. It could be expressed in some other language or by other English terms. It could be so expressed even if the word 'vixen' never existed.

There is another way to see limitations on what we can learn merely from definitions. Suppose that, although 'vixen' had always meant the same thing as 'female fox', *both* terms had meant something else, for example 'wily creature'. In that case, 'All vixens are female' would still have expressed an analytic truth, but not the one it now does. It would have meant what we now mean by 'All wily creatures are wily creatures'.

Moreover, although one can come to know that all vixens are female *through* understanding definitions of terms that now express this truth, one cannot know it wholly on the *basis* of the truth of those definitions. A route to a foundation is not itself a foundation.[1] To know that all vixens are female by virtue of knowing that, say, 'vixen' has the same meaning as 'female fox', we need a bridge between knowledge of linguistic convention and knowledge of vixens. Consider one thing such a bridge requires. We must be justified in believing a general principle something like this: that a proposition expressed by a subject–predicate sentence such as 'All vixens are female' is true if its predicate term—here 'female'—expresses something contained in the concept designated by its subject term, here 'vixen'. But this bridge principle is a good candidate for an analytic truth. If it is analytic, then, on pain of generating an infinite regress, one can know an analytic truth by knowing conventions only if one *assumes* some other analytic truth.

Moreover, to know, in the light of this bridge principle, that all vixens are female, we must take the relevant sentence, 'All vixens are female', to be the kind of thing the principle applies to, that is, to be a sentence with a predicate that expresses something contained in the concept designated by its subject. We are in effect using logic as well as knowledge of meaning to discern something about a particular sentence and to bring that sentence under a generalization about sentences. But how can conventionalism account for our knowledge (or justified belief) of the logical truths we thereby depend on, such as that if all sentences of a certain kind express truths, and this sentence is of that kind, then it expresses a truth?

The conventionalist cannot respond by doing the same thing all over again with this logical truth; for that would presuppose logic in the same way, and the procedure would have to be repeated. The problem would arise yet again. No finite number of steps would explain our justification, and an infinite number would not be possible for us, even if it would help. We could thus never account for knowledge of a given logical truth without presupposing knowledge of one. Since conventionalism presupposes (at least) logical truths of reason, in order even to begin to account for analytic ones, it cannot show—and provides no good reason to believe—either that every truth of reason is grounded in convention or even that all knowledge of such truths is grounded in convention.

Conventions as grounds for interpretation

These criticisms should not be allowed to obscure a correct point that emerges from reflecting on conventionalism. The meaning of 'vixen' *is* crucial for what proposition is expressed by the sentence 'All vixens are female', that is, for what one is asserting when (in the normal way) one uses this sentence to make an assertion. Thus, if 'vixen' came to mean the same as 'wily creature', that sentence would express a falsehood, since there are plenty of wily males. But from the fact that change in what our terms mean can result in our saying different things in uttering the same words, nothing at all follows regarding whether *what* we say in using these words is necessarily true, or true at all. Those matters depend on what it *is* that we say.

There are, however, insights underlying conventionalism: truths of reason are associated with meanings; they can be known when meanings are adequately understood; and they can sometimes be shown through pointing out relations of meanings. Moreover, without conventions, our "words" could not be said to have meanings: strictly speaking, we would have no words and could not plausibly call anything true by virtue of meaning.

Important as these points about conventions are, they do not support the conventionalist view that the truths of reason themselves, or even our justification or knowledge regarding those a priori propositions, are *based* on what words mean or on our conventions for using them. For all that these points establish, our understanding of word meanings (including sentence meanings) is simply a route to our grasping of concepts and shows what it does about the truths of reason only because of that fact.

Some difficulties and strengths of the classical view

Of the accounts just considered, then, the classical view of the truths of reason and our knowledge of them apparently stands up best. But there are other accounts and many variants on the ones discussed here. Moreover, I

have sketched only the main lines of the classical view and only some of the challenges to it. There are still other difficulties for it.

Vagueness

Recall the problem of vagueness. Perhaps the concept of being red, as well as the term 'red', is vague. Is it, then, an a priori truth that nothing is red and (any shade of) orange all over? And how can we tell?

One answer is that although words are by and large vague, concepts are not, and what *is* red (i.e., what instantiates the concept of redness) is never orange even though we have no non-arbitrary way of precisely specifying the limits of colors. Thus, we might confront a sentence, say 'That painting has a patch that is at once red and orange', which we cannot assess until we see whether it implies the necessary falsehood that the patch is two different colors all over at once or, because of the vagueness of its terms, expresses (say) the possible truth that the patch has a single color that can be considered red just as appropriately as orange.

This answer is only the beginning of a solution to the problem of how to deal with vagueness and is less plausible for highly complex concepts such as that of a work of art. The more vague our terms, the harder it is to discern what propositions are expressed by sentences using those terms, and thus the harder it is to decide whether these sentences express truths of reason. None of this implies, however, that there are *no* clear cases of synthetic a priori truths. Perhaps the proposition that nothing is round and square, taken to belong to pure geometry, is an example. (There may also be examples in the moral domain, an important possibility considered in Chapter 12.)

Meaning change and falsification

A related problem for the classical view emerges when we consider the close connection (which some regard as an equivalence) between what a term means and the concept it expresses. With this connection in mind, notice too that meaning can change gradually, as when we discover things about vixens a little at a time and thereby almost imperceptibly come to mean something different by 'vixen'. A point may then come at which it is unclear whether the term 'vixen' expresses the concept it now does or not and, correspondingly, whether or not what is then expressed by 'All vixens are female' is analytic.

This unclarity about what concept 'vixen' expresses need not give us reason to doubt, regarding the proposition which that sentence now expresses, that it is analytic; but it does show that it may be difficult to decide whether or not an utterance or sentence we have before us expresses an analytic proposition. That difficulty may drastically limit the usefulness of the notion of the analytic in understanding philosophical and other problems.

It might be argued, moreover, that on reflection the distinction between meaning change (semantic change) of the kind illustrated and falsification

of the proposition we started with does not hold. This point is likely to be pressed by those who think that the basic epistemological standard, the fundamental standard for judging whether a belief is justified or constitutes knowledge, is what is required for an overall account of experience. This broad standard is compatible both with many versions of empiricism and with some versions of rationalism.

To understand the difference between meaning change in a sentence and falsification of what the sentence is used to assert, it is helpful to contrast two cases. Compare (1) scientists' discovering that despite appearances vixens have such significant male characteristics that they are *not* really female—an outcome the classical theory says is, on the face of it, impossible—and (2) scientists' making discoveries about vixens so startling that we come to use 'vixen' in a new sense, one such that, although scientists deny that 'vixens' in this new sense are always female, what they are thereby saying provides no reason to doubt that what we *now* mean by 'All vixens are female' is true. Is there really a clear difference between (1) and (2)—roughly, between falsification of the belief about vixens we now hold and a change in the meaning of the terms we use to express it?[2]

Classical theorists take (2) to be possible and tend to hold that it is only because possibilities like (2) are not clearly distinguished from (1) that (1) *seems* possible. They regard the difference between (1) and (2) as clear enough to sustain their view and tend to conclude that what may seem to be a falsification of an analytic proposition is really only a change in meaning that leads us to substitute, for an analytic truth, what looks like a proposition inconsistent with it, yet is actually compatible with it. Other philosophers think that the difference is not clear at all and that future discoveries really can weigh against what the classical view calls analytic propositions.[3]

It is difficult to doubt, however, that there are *some* truths of reason, such as elementary logical principles, and such simple analytic propositions as that all vixens are female, that are both a priori and necessarily true.

Whether some truths of reason are also synthetic rather than analytic is more controversial, but it looks as if some of them are. Whether, if some of them are, those synthetic truths are also invariably necessary is also very controversial. I see no good reason to deny that they are necessary, but there may be no clearly decisive argument to show this.

If synthetic truths of reason are necessary, perhaps one must simply see that this is so by reflecting on the examples. In any case, our capacity of reason, our rational *intuition*, as it is sometimes (perhaps misleadingly) called, is a source of beliefs of simple truths of reason, such as the self-evident truth that if the spruce is taller than the maple then the latter is shorter than the former. We can know the *truth* of these *intuitively*, on the basis of understanding them rather than on the basis of premises for them or perceptual experience, even if more is required to know their *status* as, say, necessary or contingent, a priori or empirical. Moreover, reason, applied in our contemplating or reflecting on certain a priori truths, can yield both situational

justification—hence justification for holding beliefs of them—and actual justified beliefs of them. Clearly, reason can also yield knowledge of them.

The possibility of empirical necessary truth

It is one thing to say, with the classical view, that every a priori truth is necessary; the thesis that every necessary truth is a priori is less plausible. Consider the truth that sugar is soluble in water. Ordinarily this is thought to be a law of nature and as such something that must (of necessity) hold. Yet it is not self-evident and apparently not even broadly a priori: one could adequately understand it without thereby being justified in believing it, nor does it seem to follow self-evidently from anything self-evident. Indeed, it seems to be the kind of truth that can represent an empirical discovery. Proponents of the classical view would maintain that the necessity in question is not "logical" in the sense of absolutely precluding falsehood, but *nomic* (from the Greek *nomos*, for law), in roughly the sense characterizing laws of the natural world as opposed to every possible world or situation.

It does appear that we can clearly conceive of a lump of sugar's failing to dissolve in water, whereas we cannot clearly conceive of something that is (in overall shape) both round and square (if this is conceivable at all). But perhaps once the idea of solubility in water is properly qualified (in ways sketched in Chapter 12), there may no longer seem to be any more than a difference of degree between the two cases. I doubt that the difference is only one of degree, but let us leave the matter open and proceed to cases that pose a greater challenge to the classical view.

The truth that gold is malleable is arguably more basic to what gold is than solubility in water is to what sugar is. Is it even possible for something to be gold without being malleable? Compare the question whether a vixen could turn out to be male. This also seems impossible, but one difference is that whereas there are good ways of identifying specimens of gold without selecting them partly on the basis of malleability, there are no comparably good ways of identifying vixens without selecting them partly on the basis of being female. Still, even classical theorists grant that taking the proposition that gold is malleable to be necessary does not self-evidently commit one to considering it analytic. Critics of the classical view will maintain that it is not obvious that a specimen of gold could turn out to lack malleability, yet it is equally far from obvious that adequately understanding the proposition that gold is malleable is sufficient to justify it.

If we move to a theoretical identification statement, such as that water is H_2O, it seems even less likely that we have a proposition that is contingent rather than absolutely necessary, yet it also appears that the proposition is not a priori. The basis of our knowledge of it is confirmed scientific theory, not understanding. To be sure, there is "heavy water," but its existence bears on the kind of hydrogen atom, not on whether water of the everyday kind

is necessarily H_2O. In any case, a different kind of example also strongly supports this conclusion that some necessary truths are empirical. This time we turn to the domain of biology.

Essential and necessary truths

As the identity of human beings is normally understood, who they are is essentially tied to their parents. Is it possible that *I* might have had (biologically) different parents? Surely anyone otherwise like me but born of different parents is only a fortuitously identical "twin." Here, then, is an empirical proposition (that I am the son of R and E) which is apparently necessary.

Notice, however, that the proposition that I have the parents I do is singular and existential, implying the existence of the particular thing it concerns (me), whereas the clear cases of necessary truth we have considered are all general and non-existential. To say that nothing is both round and square, for instance, does not entail that there is anything round or square: it says roughly that anything which is round is non-square (and vice versa), and it would be true even if all the round and square things in the universe had been destroyed (and presumably even if there never had been any except perhaps in the mind of someone contemplating creating them).

What a proponent of the classical view might say of the parentage case is that the proposition that I have the parents I do is an *essential truth*—one attributing to a thing a property absolutely essential to it, roughly in the sense that it could not exist without it—but not a necessary truth. The idea is roughly this: a necessary truth holds in any possible world or situation; an essential truth holds in, but only in, those possible worlds or situations in which what it is about exists.[4]

One trouble with this view is that even in a world without water, we could speak of water and H_2O as we can of what is round or square. Perhaps the best the classical view can do here is, first, to distinguish between two kinds of necessary truth, those applicable to entities that must exist, such as (arguably) numbers, and those applicable to entities that need not exist, and second, to argue that the former truths are a priori. The idea might be that necessary truths are grounded in the nature of things, and that the nature of the kinds of things that must exist is knowable through the use of reason. The nature of water must be discovered by scientific inquiry; that of the abstract property of roundness is apparent to adequate reflection.

The idea that necessary truths are grounded in the nature of (the relevant) things has some plausibility. At best, however, it does not in any obvious way apply to purely formal necessary truths, such as that if some *A*s are *B*s, then some *B*s are *A*s, where *A* and *B* are variables and do not stand for anything in particular (they figure in indicating the form of the truth in question but provide no content).

Necessity, apriority, and provability

There is, moreover, a further objection to extending the idea to imply the apriority of all necessary truths. A theorem (in one sense of the term) might follow from a necessarily true proposition and thereby be a necessary truth—as what follows from a necessary truth is itself necessarily true—yet not be a priori because there is no way to know it simply through adequately understanding it or through adequately understanding its entailment by self-evident steps from something that is self-evident. We must not simply assume that every such theorem is self-evidently entailed by a *self-evident* proposition, or that some proof of it must proceed by *self-evident* steps from a self-evident proposition. This assumption is far from obvious and not self-evident, and the classical view must establish it by argument. It is not clear that a cogent one can be found.

It should be stressed, however, that although a provable proposition need not be self-evident, a self-evident proposition may be provable. Self-evident propositions are knowable *without* proof, on the basis of adequately understanding them, and hence are not, as are many theorems, *premise-dependent*. But many can be proved, and some may need proof in order to be accepted by some people.[5]

Moreover, even apart from those points, the only possible proof by self-evident steps from a self-evident axiom might be long; this would put the theorem a long inferential distance from the self-evident axiom(s). Granted, such a theorem would still be *provable* from what is self-evident. But simply being thus provable (yet not self-evident) entails only being what I call ultimately a priori. That status is consistent with the possibility that, for finite minds, knowledge of the proposition depends on memory. The status is thus not sufficient for an uncontroversial kind of apriority.

It appears, then, that there can be necessary truths knowable only through the work of empirical investigation or of arduous mathematical proof of a kind that cannot ground what we might call strictly a priori knowledge. Those truths, to be sure, might be both provable and knowable just on the basis of *a* use of reason—though knowledge based on a long proof also seems to depend on memory. Not just any use of reason, however, qualifies knowledge reached through it as a priori.

From the falsity of the classical thesis that every necessary truth is a priori, it does not follow, of course, that the classical view is mistaken in positing synthetic a priori knowledge or in claiming that every a priori proposition is necessary. (See Figure 6.1 for a brief representation of the classical and revised views of the a priori.)

Reason, experience, and a priori justification

Reason—conceived roughly as our mental capacity of understanding, especially in conceptual reflection or in inference—is a basic source of belief, justification, and knowledge. Like introspective consciousness and unlike

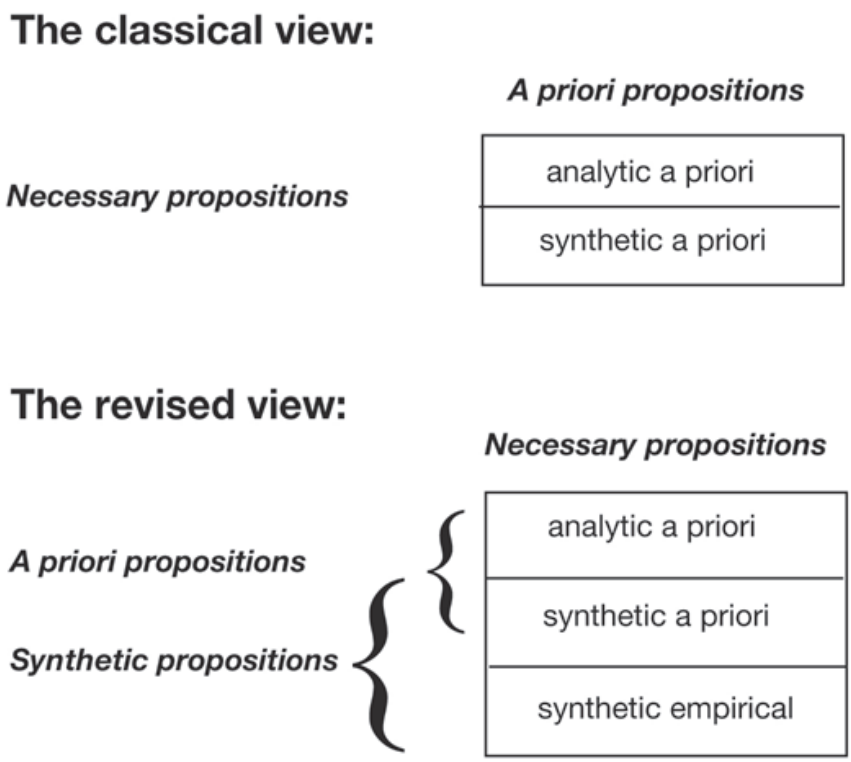

Figure 6.1 The a priori, the analytic, and the necessary.

perception and memory, it is an *active* capacity, in that we can, within limits, employ it successfully at will. I can, simply because I want to, reflect on logical and mathematical propositions. But although I can look around me just because I want to, whether I perceive anything depends on there being something there: trees and roses and books are not available to the eye in the same unfailing way that concepts and numbers are available to the mind. Through reflection on the huge range of objects of thought, we can acquire a vast amount of justified belief and significant knowledge.

To maintain that there is a priori knowledge and justification does not commit one to denying that reason has a genetic dependence on experience. Reason yields no knowledge or justified belief until experience, whether perceptual, reflective, or introspective, acquaints us with (or develops in us) concepts sufficient for grasping a priori propositions. But despite this genetic dependence of reason on experience, in one way reason may be an even firmer basis of justification and knowledge than experience. If experience is the ground from which reason grows, it is not the sole determinant of the range or power of reason. The view from the top of the tree may be more comprehensive than the view on the ground.

A priori beliefs

The notion of the a priori is not commonly applied to beliefs, but it should be clear from what has been said not only that it has a significant application to them but also that apriority on the part of a belief tends to indicate

some degree of justification. The following plausible *principle of justification for a priori belief* is a partial indication of the justificatory power of reason: *normally, if a rational person believes a proposition solely on the basis of* (*adequately*) *understanding it—believes it in a strictly a priori way—this belief is prima facie justified.*[6] In the typical cases in which this applies, the proposition, upon comprehending consideration by a rational person, will intuitively seem to the person to be true. Such an *intuitive seeming*—which for some philosophers is the primary element designated by 'intuition'—is a source of prima facie justification. We may leave open whether this, rather than the understanding in question, is the main source of the person's justification when the proposition in question is not self-evident. Plainly, however, the intuitive seeming presupposes at least a minimally adequate understanding of the proposition.[7]

There is a counterpart plausible epistemic principle—call it *a principle of knowledge for correct a priori beliefs*—to the effect that normally, if a rational person believes a true proposition in the a priori way just described, this belief constitutes knowledge. Believing in this a priori way is appropriate to (and typical for) beliefs of a priori propositions (though they may also be believed on the quite different basis of testimony), but it does not entail that the object even of a true a priori belief is a priori or a necessary truth.

It may also be true that normally, if one believes a proposition solely on the basis of one or more premises that self-evidently entail it and are themselves believed in the a priori way just described, this belief is prima facie justified. Again, such a proposition need not be a priori, but this principle is highly appropriate to what is a priori in the broad or the ultimate sense—not self-evident but either self-evidently entailed by something that is, or provable by self-evident steps from a self-evident proposition. What the principle expresses is the idea that normally self-evident entailment transmits the kind of justification that is based solely on understanding: specifically it carries that justification across a self-evident entailment. Hence, normally, if you believe a proposition on the basis of believing, with this kind of justification, a second one which self-evidently entails the first, then your belief of the first is also justified.

If these principles seem too permissive, note that we do not normally believe propositions in the strictly a priori way in question unless they are a priori and thus *can* be known on the basis of understanding them. We normally have no tendency whatever to believe, solely on the basis of understanding them, propositions indicating the state of the weather or describing the objects in our environment or the well-being or plans of others. Philosophers commonly say of such propositions that we cannot "determine a priori" (or tell or know a priori) whether they are true, and here 'a priori' designates an a priori way of believing rather than the status of the propositions in question. Compare how much we believe on the basis of perception, memory, and introspection; not only is this far more than is normally believed on the

basis of conceptual understanding, it is also quite different in the kind of grounding of the resulting beliefs.[8]

Loose and strict senses of 'a priori justification' and 'a priori knowledge'

So far, I have been speaking of knowledge and justification arising from believing in a strictly a priori *way*. This is not necessarily a priori knowledge or a priori justification, just as not everything perceptually believed is perceptual knowledge or perceptually justified. When knowledge or justification that arises from believing in an a priori way is not strictly speaking a priori, one might still call it a priori knowledge or a priori justification in the loose sense. Let us consider justification first.

Consider the proposition that people tend to feel offended when they are insulted. This is vague, but not too vague to enable us to see that it is not an a priori truth (it seems empirically true or false, since it concerns what psychological reaction a kind of conduct in fact tends to elicit). Still, imagine someone who thinks that insulting someone self-evidently entails being offensive to the person and that feeling offended is necessarily appropriate to what is offensive and tends to occur when one is insulted. Such a person might argue that, on the basis of understanding it, we can believe the proposition that people tend to feel offended when insulted, and that we may, on this basis, be justified in believing that. If one might be so justified, then we might speak of a priori justification in the loose sense. We may also say that the belief itself is a priori in the loose sense, since it is grounded in an a priori way: if it is not grounded in the strictly a priori way (based solely on an adequate understanding of the proposition), the belief is at least held in *an* a priori way—it is based solely on an understanding of the proposition. Just as a perceptual belief can be justified and false (as when one first sees a straight stick half submerged in water and thinks it is bent), this belief can be also.

Another case of a priori justification in the loose sense can occur when, although one believes a proposition that *is* a priori, one believes it on the basis of an *inadequate* understanding of it. This is still believing it in an a priori way, however, as the basis of one's belief is one's understanding of the content of the proposition. But it is not believing in a strictly a priori way, as that requires adequate understanding. One might, for instance, overlook a subtlety or confuse one notion with a similar one, such as believing a proposition and being disposed to believe it. Suppose that, on the basis of my understanding of it, I believe a mathematical theorem that is a priori in the broad sense. Suppose further that this understanding, although inadequate, is not unreasonable (say because it represents a plausible though subtly misguided interpretation of the theorem). Then my belief may be justified. This is a second case of a belief held in an a priori *way* and exhibiting a priori justification in the loose sense. Here the proposition *is* a priori, but the justification,

though based on a reasonable understanding, is defectively grounded. In the other case of a priori justification in the loose sense, the belief is also held in an a priori way, but the proposition is not a priori.

If a belief that is a priori justified in the loose sense could constitute knowledge, we might speak of a priori knowledge in the loose sense. But as both our examples of such justification exhibit a *defective* (though reasonable) understanding in the basis of the justification, they are not plausibly considered instances of knowledge. Beliefs resting on a basis embodying conceptual error are not plausibly taken to constitute knowledge, even if the conceptual error is justified.

Suppose, however, that I believe a mathematical theorem on the twofold basis of a self-evident axiom (which I adequately understand) and the justified true belief that the theorem is entailed by the axiom (we may assume the second belief to be grounded wholly in my mathematical knowledge and understanding). Suppose further that the theorem is entailed, but not *self-evidently* entailed or self-evident.[9] It is not self-evidently entailed because adequately understanding the conditional proposition that if the axiom holds then the theorem does is not sufficient to justify believing this conditional. To see the truth of this conditional proposition, I must note several intermediate steps from the axiom to the theorem, so that I do not see its truth (or the entailment it expresses) on the basis of adequately understanding the proposition. Still, the entailment is provable, and by proving it I may know the theorem. This is surely a broadly a priori way of knowing it, and the proposition itself is, in my terminology, ultimately a priori. Correspondingly, we may speak of a priori knowledge in the loose sense here. The knowledge is not a priori in the *strict* sense because the theorem is not a priori, even in the indirect sense. By valid deduction, I can prove it using the a priori procedures illustrated, but such provability of a proposition is not sufficient for its being self-evident or even knowable a priori in the strict sense of that phrase.

By contrast, a priori knowledge in the strict sense is not only more than true belief held in a strictly a priori way; it is also more than knowledge of an a priori proposition. I could know a simple logical truth on the basis of testimony, even if it *can* be known on the basis of understanding alone. This would be knowledge of an a priori proposition that is not even a priori knowledge in the loose sense. Its grounding (wholly) in testimony does not prevent its being knowledge, but testimonial grounding of a belief does preclude its constituting a priori knowledge of any sort. Again, the analogy to perception is helpful. Just as perceptual knowledge is knowledge based on perception and thus more than knowledge about a perceptible, a priori knowledge is knowledge based on understanding and thus more than knowledge of an a priori proposition.

To achieve a more specific characterization of a priori knowledge we do well to begin with a crucial constituent of it—*a priori justification*. In the strict sense (the sense that mainly concerns us), this is justification based

directly or indirectly on understanding a self-evident proposition (the justification will be only situational if the person in question does not believe the proposition). A priori justification (in the strict sense) thus divides into two kinds, depending on whether it is directly or indirectly based on understanding some self-evident proposition. (1) A priori justification for believing a proposition is based *directly* on such understanding when the justification depends only on understanding that proposition itself. This is a priori justification in the strict and narrow sense. (2) A priori justification for believing a proposition is based *indirectly* on such understanding when the justification depends on *also* understanding a self-evident entailment of that proposition by some self-evident proposition. This is a priori justification in the strict but broad sense.[10]

If this outline is correct then *a priori knowledge*, in the strict sense, might be plausibly taken to be knowledge that is based, directly or indirectly, in the way just indicated, on understanding one or more self-evident propositions. There is, then, in addition to a division between a priori justification and a priori knowledge in the strict and loose senses, a division between direct and indirect (non-inferential and inferential) a priori justification, and direct and indirect a priori knowledge, in both senses.[11] (Figure 6.2 represents the four dimensions of the a priori we have been exploring.)

The power of reason and the possibility of indefeasible justification

We have seen that, and perhaps to some extent how, the justificatory and epistemic power of reason enables it to ground a priori knowledge and a priori justified beliefs of a priori propositions. We have also seen its power to provide such knowledge and justification, in loose senses of 'a priori knowledge' and 'a priori justification', for propositions that are not a priori but invite belief on the basis of their conceptual content. These senses are especially appropriate for propositions that are provable from what is a priori. Is the power of reason such that it provides for something that even introspective experience apparently does not—indefeasible justification? It will help to focus on a concrete example.

There may be truths of reason that are so simple and luminously self-evident that they *cannot* be unjustifiably believed, at least at a time when one comprehendingly considers them. Could one comprehendingly consider, yet unjustifiably believe, that if Shakespeare is identical with the author of *Hamlet* then the author of *Hamlet* is identical with Shakespeare? This is doubtful. One could perhaps believe it partly on the basis of a bad argument; if one did, there would be something unjustified in the *way* one believes it. But if one believes it, one has some understanding of it, and if one understands something this simple to the extent required for believing it, it is at best difficult to see how one could fail to have an understanding of it adequate to yield

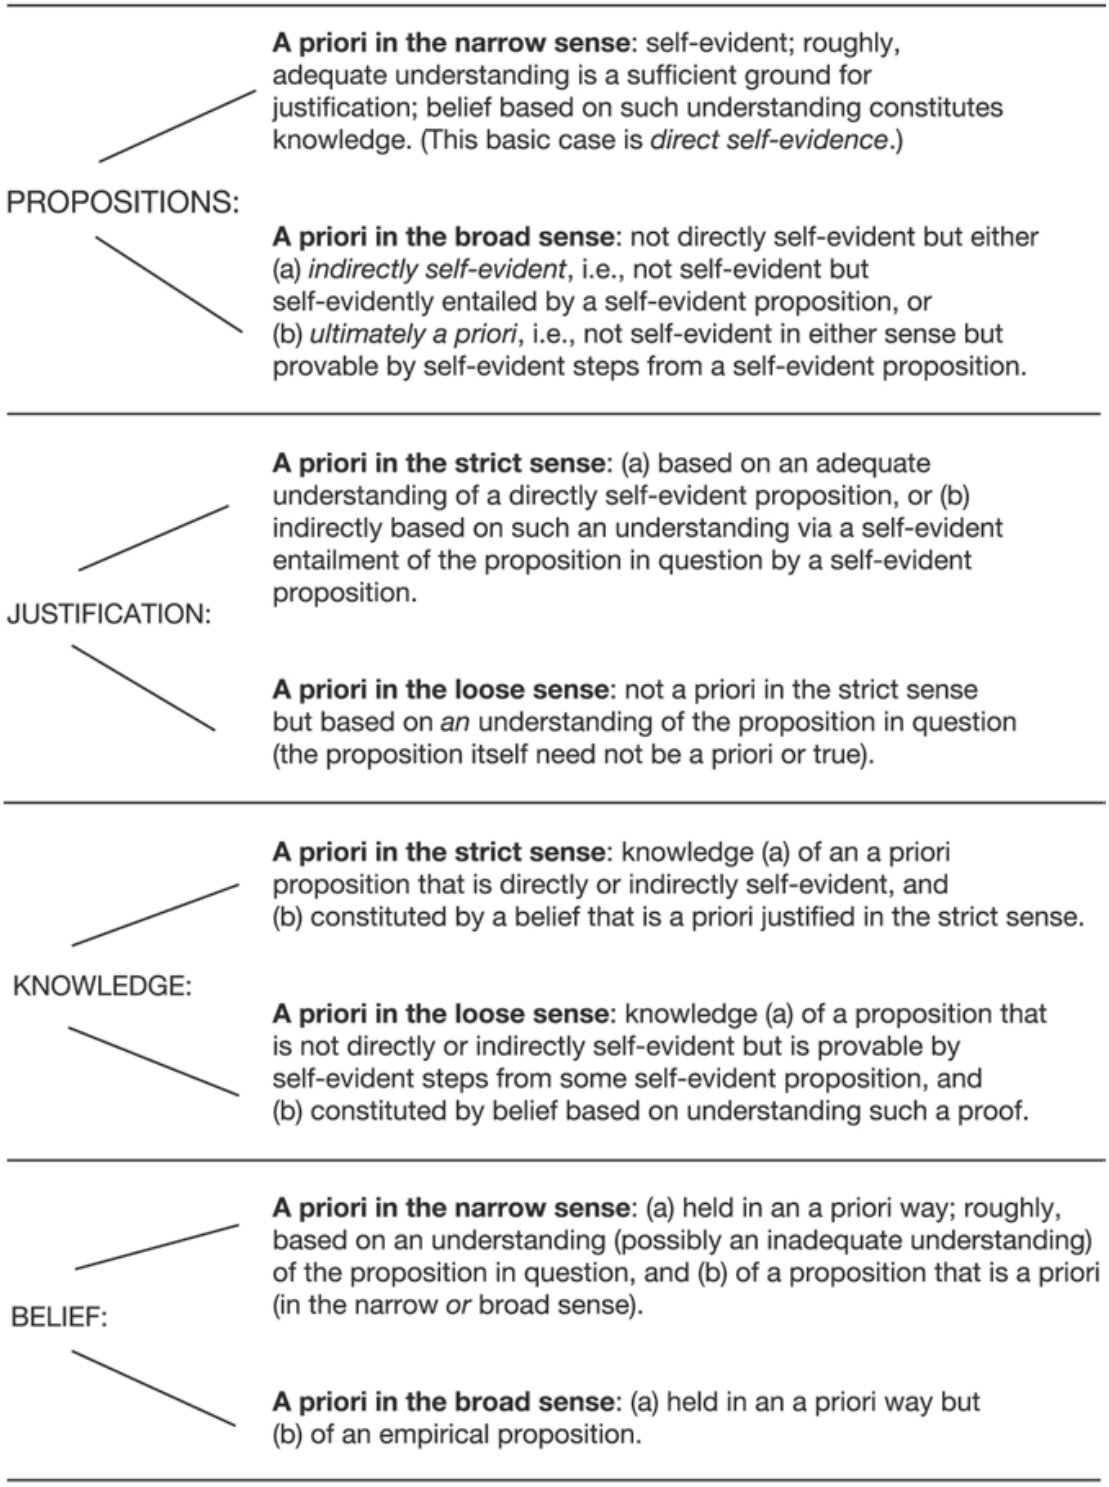

Figure 6.2 Outline of a four-dimensional conception of the a priori.

justified belief of it, at least at a time when one comprehendingly considers it. Perhaps, then, a belief held under these conditions would be—or at least could be—indefeasibly justified.

If there are propositions like this then there can apparently be indefeasible justification: justification so secure that those possessing it cannot be

unjustified in believing the proposition in question.[12] But not all a priori justification (even in the strict sense) should be considered indefeasible. Justification for believing even certain logical truths can be defeated by plausible skeptical arguments.

Perhaps, moreover, not all presumptively indefeasible justification need be a priori. Consider my justification for believing that I exist, a proposition that is neither a priori nor necessary but is arguably such that I cannot unjustifiably believe it. If there is indefeasible justification, this is important in dealing with skepticism (as Chapter 13 will), but plainly such justification is not a characteristic mark of either a priori or empirical justification. If, on the other hand, there is no indefeasible justification (something I leave open here), at least our understanding of simple self-evident truths of reason gives us both very secure justification for believing those truths and, when we do believe them on the basis of adequately understanding them, knowledge of them.

In summarizing some apparently warranted conclusions regarding the truths of reason, we might focus on how much seems plausible in the classical view that the a priori is coextensive with the necessary but includes the analytic as a subcategory: that any proposition that is a priori is necessary and conversely, but not every a priori proposition is analytic. Apparently, it is true that not all propositions knowable on the basis of adequately understanding them are analytic. The classical view seems correct in its claim that not everything a priori is analytic. It seems mistaken, however, in the idea that every necessary proposition is a priori, though probably not in the plausible idea that every a priori proposition is necessary.

More positively, in addition to our having a priori knowledge of self-evident propositions, on the basis of such knowledge we may know many truths that are at least ultimately a priori: not themselves self-evident but self-evidently entailed by, or provable by self-evident steps from, some proposition that is. Many of our beliefs, most clearly certain logical and mathematical ones, are grounded in understanding of their content. Reason, then, as manifested in our capacity for understanding, is one of the basic sources of belief, justification, and knowledge; and, in a way that the other three sources we have explored do not, it enables us to know truths that hold not only in the world of our experience but also in any circumstances whatever.

Notes

1 At least in his classic 'Two Dogmas of Empiricism', in his *From a Logical Point of View* (Cambridge, MA: Harvard University Press, 1961), W.V. Quine sometimes talks as if he thinks that a knowledge of synonymy (sameness of meaning) of words is necessary for any possible knowledge of analytic propositions. See, for example, section 4, on semantical rules. One important comment is that "definition turned out to be a

will-o-the-wisp, and synonymy turned out to be best understood only by dint of a prior appeal to analyticity." In the overall context, the suggestion may be that only an independent conception of synonymy would clarify analyticity.

2 Cf. W.V. Quine's remark that "truth in general depends on both language and extra-linguistic fact. The statement 'Brutus killed Caesar' would be false if the world had been different in certain ways, but it would also be false if the word 'killed' happened rather to have had the sense of 'begat' " ('Two Dogmas', section 4). Compare saying that the *sentence* 'Brutus killed Caesar' would have expressed a different, and false, proposition (which is what defenders of the classical view would likely say). Has Quine provided any reason to think that the *statement* in question—understood as the historical truth we express using the sentence—would have been false if the English word 'killed' had meant 'begat'?

3 For a valuable discussion of the notion of the analytic in relation to the conceptual, see M. Giaquinto, 'Non-Analytic Conceptual Knowledge', *Mind* 105, 418 (1996), 249–68. One of his major conclusions bears on the status of such cases as the proposition that all vixens are female:

> What the liberated position [Quine's, freed of behaviorism] maintains is that any belief may be rationally rejected in the light of future findings; what it has to accommodate is that some beliefs may be rationally retained even when their customary linguistic expressions become unacceptable. These [positions] are not inconsistent.
>
> (p. 266)

4 The terminology of possible worlds traces especially to Gottfried Wilhelm Leibniz and has been influentially discussed in relation to a number of the issues concerning necessity and the a priori by Saul Kripke in *Naming and Necessity* (Cambridge, MA: Harvard University Press, 1980). Kripke offers a different kind of example of empirical necessities: true identity statements formed using proper names, as in 'Hesperus is identical with Phosphorus' (both being names of Venus). He also argues, using the example of the standard meter stick in Paris, that an a priori truth, say that the length of the standard meter stick in Paris at time *t* is 1 meter, may not be necessary. This is a highly controversial example (more often attacked than defended), which I cannot discuss here. For detailed criticism, see Albert Casullo, 'Kripke on the A Priori and the Necessary', *Analysis* 37 (1977), 152–9. Casullo also usefully distinguishes knowledge of the *truth value* (truth or falsity) of a proposition from knowledge of its *modal status* (its being necessarily true or false, or contingently true or false), and argues that the classical view could be mistaken in holding that the truth value of necessary propositions is

always knowable a priori yet correct in holding that their modal status is knowable a priori.

5 Many philosophers have taken self-evident propositions to be unprovable, e.g. W.D. Ross (*The Right and the Good*, Chapter 2), apparently following G.E. Moore and others. A simple counter-example is the proposition that if *p* entails *q* and *q* entails *r*, then *p* entails *r*.

6 Two comments are needed here. First, it might be desirable to widen the characterization to allow beliefs based *at least predominantly* on understanding the proposition in question (which requires understanding the concepts figuring in the proposition); but I want to avoid here the complications that arise from considering multiple bases; thus I shall not generally qualify 'based on' and similar terms. The main points in question will hold if it is taken as equivalent to 'essentially based on'. Second, although the relevant beliefs might be thought to be *always* prima facie justified, there is at least one difficulty with this: perhaps there could be an abnormal case of a kind that prevents *any* justification from arising. This is not obviously possible, since if understanding is a sufficient basis for the belief, that might arguably carry some degree of justification. In any case, the normality formulation is significantly strong.

7 The view that phenomenal seemings (including perceptual as well as intuitive seemings) suffice for justification is commonly called *phenomenal conservativism.* The position is defended by, e.g., Michael Huemer in *Skepticism and the Veil of Perception* (Lanham, MD: Rowman and Littlefield, 2001). For critical discussion of the view see Matthias Steup, 'Internalist Reliabilism', *Philosophical Issues*, 14 (2004), 403–24.

8 The quantitative comparison may be challenged by those who think we have infinite sets of mathematical beliefs (e.g. that 2 is even, 4 is even, etc.) and beliefs based on others by trivial operations, such as forming new beliefs by adding an 'or', as when, given my belief that I am seated, I form the belief that either I am seated or I am flying to the moon. That this conception of belief is mistaken will be argued in Chapter 9, which also notes relevant literature. In any case, the contrast I am drawing here would be adequately strong even without its quantitative dimension.

9 As indicated in explicating self-evidence, self-evident entailment (as opposed to entailment simpliciter) is not transitive. If it were then if an axiom, *A*, self-evidently entailed a theorem, which self-evidently entailed another, and this held for 100 steps to theorem *T,* the proposition that if *A,* then *T* would have to be self-evident. But reflection on axiomatic systems shows that this is not so.

10 This implies that even if one justifiedly believed, and knew, an a priori proposition on the basis of a self-evident axiom, but *not* on the basis of a self-evident entailment of the former by the latter (say, by a chain of non-self-evident but valid inferences instead), the justification and knowledge would still not be a priori in the strict sense—though they might be very close to it.

11 Three comments are needed here. First, for one's justification to be a priori, at least in the strict sense, it must not depend (epistemically) on memory. Thus, suppose there are too many self-evident premises for me to hold in mind at the same time as I understand the proposition that my conclusion follows from them. Or, suppose there are so many self-evident steps linking a single self-evident premise to a conclusion that I cannot hold them all in mind in a way that assures understanding the ultimate entailment of that conclusion by the premise. Then *my* justification for believing this conclusion is not a priori (though I may be able to prove the conclusion). Second, and related to this, so long as there can be a mind sufficiently capacious to understand the entire set of propositions in question (the premises and the proposition that if they are true, then the conclusion is also) without dependence on memory, a priori justification for *someone's* believing the conclusion is possible. Third, as in this book generally, I regard the justification referred to as defeasible (a notion considered in this chapter and again in Chapter 11) unless otherwise specified.

12 It might be argued, however, that if one believed such a simple self-evident proposition *essentially* on the basis of a bad argument, one would not *justifiedly* believe it, though, by virtue of adequately understanding it, one would still *have* a justification for believing it which simply fails to serve as a sufficient ground of one's belief. I leave open whether one could believe such a proposition both fully comprehendingly and essentially on the basis of a bad argument (as opposed to one's being only influenced by such an argument).

7 Testimony

The social foundation of knowledge

- **The nature of testimony: formal and informal**

- **The psychology of testimony**
  The inferentialist view of testimony
  Inferential grounds versus constraints on belief-formation
  The direct source view of testimony
  Testimony as a source of basic belief

- **The epistemology of testimony**
  Knowledge and justification as products of testimony
  Testimony and memory compared
  The twofold epistemic dependence of testimony

- **The indispensability of testimonial grounds**
  Conceptual versus propositional learning
  Testimony as a primeval source of knowledge and justification
  Non-testimonial support for testimony-based beliefs

7 Testimony

The social foundation of knowledge

If our only sources of knowledge and justified belief were perception, consciousness, memory, and reason, we would be at best impoverished. We do not even learn to speak or think without the help of others, and much of what we know depends on what they tell us. Children in their first years of life depend almost entirely on others for their knowledge of the world.

If perception, memory, consciousness, and reason are our primary individual sources of knowledge and justification, testimony from others is our primary social source of them. This is why it is a primary concern of social epistemology. The distinctive situations in which testimony yields knowledge and justification are social: in each case one or more persons convey something to one or more others. There are various kinds of testimony, however, and there are many questions about how one or another kind yields knowledge or justification.

The nature of testimony: formal and informal

The word 'testimony' commonly evokes images of the courtroom, where formal testimony is given. Someone sworn in testifies, offering information supposed to represent what the person knows or believes. Often such testimony recounts what was witnessed first-hand, but testimony can be an expression of what we believe about something we did not witness, such as the implications of a scientific theory or the potentials of human character.[1]

Formal testimony differs from the informal kind in the conditions under which it is given, but not necessarily in being more credible. Testimony of the informal kind—roughly, saying something in an apparent attempt to convey (correct) information to someone else—plays a very large role in our lives and raises the question of the importance of testimony for knowledge and justification.[2]

For the informal giving of information, for instance in telling someone where one was last night, 'testimony' is too heavy a word. We could speak of 'informing', but this is too narrow, both in suggesting a prepared message (as in 'Yesterday she informed me of her plan to attend') and in (normally)

implying that what is conveyed is true. We might regard all testimony as a kind of saying. But not all saying—even apart from what is said in fiction—is testimony. Someone who says, 'Ah, what a magnificent tree!' is expressing a sense of the magnificence of the tree, but not giving testimony that it is magnificent, as when an arborist cites features of shape and color in supporting a claim that the tree is magnificent and worth the high cost of pruning.

For much conveyance of information it can help to speak of *attesting*. This covers both formally testifying that something is so and simply saying, in the relevant informational way, that it is so, for instance telling someone the time. Testimony is always given to one or more persons (to oneself, perhaps, in the limiting case). It may be actual or, in some cases, hypothetical, as when a diarist describing atrocities for posterity does not know whether anyone will read the testimony. In any event, what we must understand here is the role of testimony of all these kinds—roughly, of people's telling us things—in accounting for our knowledge and justification. I begin with the psychological question of how testimony yields belief. The psychology of testimony is both intrinsically interesting and epistemologically important.

The psychology of testimony

If we start thinking about testimony by focusing on formal cases, we might conclude that as a source of belief, testimony is quite unlike perception in that testimony produces in us only inferential beliefs of what is said, whereas perception produces non-inferential beliefs about what is perceived. The idea that beliefs based on testimony arise by inference from one or more premises is probably a natural result of concentration on formal testimony. When I hear courtroom testimony, I appraise the witness, place the testimony in the context of the trial and my general knowledge, and accept what is said only if, on the basis of this broad perspective, it seems true. I do not just believe what I hear, as I may just believe that a bat flew by if I see one zigzag across the evening sky. Sometimes it is like this: given the premises that (for example) the witness seems credible and that the statement in question—say that the accused dined in a certain restaurant on New Year's Eve—fits what I know, I may thereby come to believe this statement. Let us assess the idea that testimony-based beliefs in general arise in this inferential way.

The inferentialist view of testimony

If this inferentialist picture of testimony is correct, then testimony is a significantly less direct source of belief than perception: it yields belief only through both the testimony itself and one or more premises that support the proposition attested to or the attester's credibility. If that is so, testimony is also not as direct a source of knowledge or justification; for one would know, or be justified in believing, what is attested only if one knows, or is at

least justified in believing, one's premise(s). One could not know simply *from* testimony, but only from premises *about* it as well.

There is a different, and I think more plausible, account that can also explain the psychological role of certain background beliefs. On this account, beliefs about the credibility of the attester and beliefs pertinent to the attested proposition may play a mainly filtering role. These beliefs (among other filters) prevent our believing testimony that does not "pass," for instance because it seems insincere. But if no such difficulty strikes us, we "just believe" (non-inferentially) what is attested. These filtering beliefs (which are common but not necessarily elements in the mind of everyone who forms beliefs from testimony) are like a trapdoor that shuts only if triggered. Its normal position is open, but it stays in readiness to block what should not enter.[3]

The open position of our natural testimonial filter is a kind of *trust*. Trust is indeed apparently the evolutionary default position: without a significant degree of trust in others, children could not, without extraordinary luck, reach adulthood and our species would likely not have survived. The absence or laxity of filtering beliefs yields *credulity*; the presence of excessively rigorous ones yields *skepticism*. Intellectual virtue—and *epistemic responsibility* conceived as a kind of virtue—are attained when we achieve a reasonable "mean" between excessive credulity and unwarranted skepticism.

It could very well turn out that, in different circumstances, each of these accounts—the inferentialist account and the non-inferential filtering account—applies to the formation of beliefs of what we are told. The psychological possibilities here are numerous, and it should be stressed that beliefs are not the only filtering elements that can non-inferentially guide formation of testimony-based beliefs. An intuitive sense of plausibility may also serve to filter. Fortunately, we need not describe all the possible filters. For now, it is enough to see that we need not consider belief properly said to be *based on testimony* to be inferential, say grounded in a further belief that the attester has spoken plausibly.

In the case of informal testimony—the most common kind—the beliefs it produces in the hearer are surely not inferential. Certainly when trusted friends speak to us on matters we have no reason to think are beyond their competence, we normally "just believe" what they tell us. Indeed, if I am trusting of people's word, then normally, when people tell me something, my belief system stands ready to be stocked. I hesitate or draw cautionary inferences only if (for instance) a would-be new belief conflicts with one or more beliefs already in my inventory. If you look vigorous and tell me you once swam the English Channel, I may readily believe you, whereas without special evidence I would not believe someone claiming to have climbed Mount Everest without using rope. For on the basis of my relevant background beliefs about climbing, I take that feat to be impossible. I have filtering-beliefs that prevent the testimony's passing into my belief system.

Inferential grounds versus constraints on belief-formation

These points about how testimony produces belief need expansion. Just as it is misleading to try to build an account of the psychology of testimony from the formal cases, it is a mistake to take a *static* view of how testimony produces belief. Our beliefs and even our belief-forming processes may change in the course of our receiving testimony. I meet someone on a plane. She tells me about a conference in which a speaker I know lost his temper. Initially, I suspend judgment about whether he did so, as the incident is of a rare kind and I do not know her. Then, as she describes the conference further, other details begin to fit together very well, and she gives information I already know, such as who was there. Soon I am listening in an accepting attitude, forming beliefs of each thing she says as fast as she proceeds. At the end, I find that I *now* believe that the speaker did in fact lose his temper.

Even at the beginning, I need not have *inferred* that I should suspend judgment on the initially unlikely statement about the speaker. Suspending judgment may be a non-inferential response to the constraints set by my independent beliefs or my sense of plausibility. Moreover, her testimony is *blocked*, but not *overridden*, by my antecedent beliefs and impressions. That is, they prevent my believing what she attests to, but do not lead me to *disbelieve it*. They do not overturn a testimonially grounded belief I formed and then gave up because of what I later came to believe, as when I discover it is inconsistent with apparent facts.

What happens is apparently this. As her narrative progresses, the constraints set by my independent beliefs relax, and, regarding each statement she makes, I form beliefs not only non-inferentially, but also even spontaneously, in the sense that any constraints that might have operated do not do so. Her statements no longer have to be tested by under the gaze of my critical scrutiny, nor are they filtered out by the more nearly automatic checking the mind routinely does when people offer information.

The most difficult thing to explain here is why, at the end, I believe the proposition on which, at the beginning, I suspended judgment. One might posit an unconscious inference, say from the general credibility of her account to the conclusion that this proposition, as an essential part of it, is true. But in what sense can an inference, as a mental process, be unconscious? This is far from clear. In any case, perhaps the cognitive *influence* of my standing beliefs, such as a newly formed belief that she is credible, need not proceed through an *inference* from them. It might be like this: even apart from my forming beliefs about her credibility, her eventually becoming, in my eyes, quite credible can in some fairly direct way produce in me a general disposition to believe her. This disposition is strengthened as she speaks with an evident credibility; and at the end it overcomes the resistance to belief which was exercised earlier by my constraining beliefs. On the subject she is addressing, I have come to *trust* her. The case shows, moreover, that trust can be retroactive as well as retrospective.

The direct source view of testimony

There are still other possibilities that support the conclusion that the inferentialist view of testimony is too narrow. Perhaps people (or some of us) have a credibility scale on which attesters acquire—commonly without our conscious attention to the matter—a place that can change, also without our conscious attention. This is an interesting empirical hypothesis that I cannot pursue, but all that is crucial here is that we see how beliefs grounded in testimony—testimony-based beliefs—can be *constrained by* other beliefs without being inferentially *based on* them and how beliefs based on testimony can be formed *later* than the attestation that is their ultimate source.

Perception, too, can produce belief after it has begun or, indirectly and with the help of memory, even after it has ceased. One may look at a shape for a long time before believing that it is a tree stump and not a stroller who stopped to gaze at the night sky. This same belief could also arise much later, from vividly recalling the image a day later when one is questioned about the scene. The connection in virtue of which a belief is based on a source need not be direct or simultaneous or a result of inference from premises.

Is the analogy with perception sufficient to warrant concluding that, like perception, testimony is a *basic source of belief*, in the sense, roughly, that it can produce belief without the cooperation of another source of belief? Consider perception. If I see a tree, this can produce in me a belief that there is a tree before me without my having a potentially belief-producing experience of any other sort, such as a separate consciousness of an image of a tree.[4] But I cannot form a testimony-based belief unless I *hear* (or otherwise perceive) the testimony. Perception is crucial for the formation of testimony-based beliefs in a way that no other belief source is crucial for the formation of perceptual beliefs.[5]

Granted, perception does not produce belief without appropriate background conditions, nor does its being a basic source of belief imply that antecedent beliefs are irrelevant. Suppose I firmly believe I am hallucinating the moon. Then, even if I actually see it, I may withhold judgment on whether it is out. A basic source does not derive its *generative* power from another source, but it need not operate in complete independence of other sources or their outputs. It can yield belief without the help of another source; but it may also cooperate with other sources in producing belief, and they may suppress some of its would-be products or may undermine the justification of some of the beliefs it does produce.

Testimony as a source of basic belief

Given that testimony-based beliefs are not inferential, and so need not be grounded on a belief that the attester is sincere or even on a belief that someone is speaking to one (though one must be at least disposed to believe

this), one may be puzzled by the point that testimony is not a basic source of belief. The puzzlement may arise from failing to see that perception itself is required for the formation of belief based on testimony, even if perceptual *belief* is not a requirement.

To be sure, I may have to be *disposed* to believe someone has said that the speaker lost his temper to acquire a testimony-based belief of this statement; but that seems to be only because I must have comprehendingly *perceived* this *being* said; it does not imply that I have formed the belief *that* it was said, just as perception of a sentence in a convincing article one is reading can produce belief of what it says without one's forming the belief that the sentence says that. There is no reason to think the mind must keep such semantic double books. It is my perception of what is said, typically my hearing or reading it, that is required for formation of a testimony-based belief of the proposition attested to. There is a sense in which I must know—having taken in—*what* was said; but this is a kind of *understanding* and does not require forming such specific beliefs as that Juan said that Jane is reliable. We could speak of *recognitional knowledge* here: a kind of knowing *what*. Such knowledge need not be expressed in beliefs.

There is also a positive point here. Testimony can be a source of *basic beliefs*, in the sense of *beliefs not based on one or more other beliefs*. The beliefs testimony evokes need not be based on premises at all, much less on premises grounded in another belief source. The kind of non-inferential belief that testimony typically produces, the kind I am calling testimony-based belief, can also be basic knowledge *if* it meets the conditions for non-inferential (hence non-premise-based) knowledge. It can certainly be basic *for* a person in the everyday sense of being central in the person's life.

A major epistemological point that the case of testimony shows here is that a basic belief—roughly, one basic in the order of one's beliefs, and so not premise-dependent—need not come from a basic source of belief: roughly, one basic in the order of cognitive sources and so not source-dependent. A belief that is not based on, and in that sense does not depend on, another belief may come from a source of beliefs that *does* depend on another source of them.

The epistemology of testimony

In the light of what has emerged about how testimony produces belief, we are now in a good position to ask two further questions. How does testimony yield knowledge and justification, and does it ever yield basic knowledge or basic justification in the way perception and reflection, for instance, apparently do? The case of knowledge is in some ways easier to deal with here than that of justification, and I will start with knowledge. As with perceptual knowledge and justification, testimony-based knowledge and justification turn out to differ.

Knowledge and justification as products of testimony

Testimony gives knowledge to its hearers only under certain conditions. If I do not know that the speaker at yesterday's conference lost his temper, then you cannot come to know it on the basis of my attesting to it.[6] This is obvious if I am mistaken and he in fact did *not* lose his temper. But suppose I make a lucky guess and am right. Then I give you correct, conjectured information which I do not know; but you are also lucky to be correct and also do not know that he lost his temper. It is a fluke that I get it right; it is even more of a fluke that you get it right, since in your case there are, in addition to the chance I have taken of making a mistake, the other liabilities you escape: of my having distorted the truth, of your having misheard me, of your adding a false detail to what you take from my testimony, and so forth.

There is a more common defect in testimony that prevents its producing knowledge in the hearer. Imagine that I do not guess at, but incautiously accept, the proposition that the speaker lost his temper, from someone I know often lies about others. Again, I lack knowledge that he lost his temper, even if this time the proposition is true; and again, you cannot know it on the basis of my testimony, which is now ill-grounded in another way.

The case with justification is quite different. Even if I am not justified in believing that the speaker lost his temper, I can be credible to you in such a way that *you* can become justified in believing this on the basis of my attesting it to you. To see this, consider the two facets of *testimonial credibility*: the *sincerity* dimension, concerning the attester's honesty, and, second, the *competence* dimension, concerning the attester's having experience or knowledge sufficient to make it at least likely that if the attester holds a belief of the proposition in question or of closely related ones, then they are true. Surely you can justifiedly regard me as credible on the topic of whether the speaker lost his temper if you have good reason to believe that I am honest, possess normal acuity and memory, and was present and reasonably attentive on the occasion.

Consider a further asymmetry: I cannot give you testimony-based knowledge that something is so without having knowledge that it is so, yet I *can* give you justification for believing this without having such justification. This asymmetry is important but can mislead. In both cases my credible report (my testimony) is your basis: of your justification for believing what I attest to and (when I know it) of your knowledge of it. But whereas I transmit to you my knowledge that the speaker lost his temper (when I know this), when I do not have justification for it I do not *transmit* to you justification for believing this—I do not have that justification to transmit. Rather, the way I attest to the proposition, together with your background justification regarding me and the circumstances, gives you this justification, independently of whether I have it. This illustrates non-transmissional grounding of justification, where testimony-based knowledge is transmissionally grounded.

Thus, in normal cases in which you credibly attest to something you

know, you do not give me justification in the way you give me knowledge. Testimony-based knowledge is received by what it is natural to call *transmission* and so is dependent on whether the attester knows the truth of the proposition in question—call it *p*. By contrast, recall the speaker who attests to *p,* which the person obviously knows, in a tense way that betrays anxiety, which *p* does not mention. It is natural to say that you here gain knowledge of the anxiety *through* the testimony, whereas you would gain knowledge that *p* on the *basis* of the testimony. In the first case, the knowledge has no essential relation to what is attested to—which has no connection with anxiety—and is not testimony-based.

Testimony that *p,* then, can convey the attester's knowledge that *p*; it can *produce* in the hearer *a* justification for believing *p*, and it can even yield knowledge (whether of *p* or of some other proposition) without that knowledge being based on it. But testimony that *p* does not convey the attester's justification for believing *p*—the attester need not even have such justification. My testimony that *p*, then, does not give recipients justification in the way it gives them knowledge.

This contrast between conveying knowledge and providing justification helps to explain the original asymmetry: if I do not know that a proposition is true, my attesting to it cannot transmit to you testimony-based knowledge that it is so (I have no knowledge to give here); but even if I am not justified in believing it, my attesting to it can give you justification for believing it, through providing the main materials for your becoming justified in believing it.[7] One might claim that this is still not testimony-based justification, but I think it can be, in the clearest sense in which there is such a thing. To see this, let us compare testimony with memory.

Testimony and memory compared

The contrast between how testimony produces knowledge and how it produces justification in the recipient is reminiscent of a contrast applicable to memory (drawn in Chapter 3). Just as we cannot know that *p* from memory unless we have *come* to know it in another way, say perceptually, we cannot know that *p* on the basis of testimony unless the attester (or someone from whom the attester comes to know it) has come to know it (at least in part) in another way; whereas we can become justified in believing *p* through memory impressions, whether or not *p* is true or known,[8] and we can become justified in believing *p* on the basis of testimony, whether or not the attester has true belief or knowledge of it or even justification for it.

With testimony-based knowledge, as with memorial knowledge, there must apparently be a certain kind of unbroken chain from the belief constituting that knowledge to a source of the knowledge in some other mode, such as seeing; but with testimony-based justification, as with memorial justification, what is essential is apparently a matter of the present epistemic situation of the subject or recipient, such as the contents of apparently

memorial consciousness and the content and justifiedness of background beliefs. Memory and testimony can each generate justification (though in different ways); but they are not generative with respect to knowledge: characteristically, the former preserves knowledge, the latter transmits it.[9]

There is another way in which justification and knowledge apparently differ in their relation to testimony. Suppose I *am* justified in believing *p*, but you have no justification of your own for believing *p* or for taking me to be credible on the topic. To vary the conference example, imagine that in passing, and without giving evidence, I say that three speakers lost their tempers, and your background information neither disconfirms nor supports this claim or my credibility in the matter. Here justification follows your lights rather than mine: my would-be contribution to justifying you in believing *p* is undermined by your lack of justification for thinking my testimony is credible or for believing *p* on some other basis. Receptivity to testimony-based justification sometimes requires already having some measure of justification: for believing the attester credible or for believing *p*, or for both.

Knowledge is different on this score: to know something through my attesting to it in expression of my own knowledge, you do not have to know that I am credible; it is quite enough that you have some reason to believe I am and no reason to doubt it. It is normally enough that you presuppose it and have no reason to doubt it. Surely you can know that it is nine o'clock on the basis of my knowing this and telling it to you, even if you simply find me a normal-seeming person with a normal-looking watch and take me to be credible.[10] And why indeed must you meet any more than a negative condition: not having reason to doubt my credibility? After all, we are talking about a case in which I know that it is nine o'clock, attest to this *from* my knowledge of it, and thereby produce your (true) belief that it is nine. There is, then, a kind of unbroken chain from the fact that it is nine to your true belief that it is.

A natural objection to this credible-unless-otherwise-indicated view of testimony as a ground for knowledge is that in our example one's evidence is so scanty that one would at best have only some reason to believe it is nine o'clock. But is this true? Granted, my having some reason to believe the proposition may be all I can *show* from my evidence or from what I feel certain of. Still, on the assumption that I in fact do know the time and sincerely tell it to you, it would seem that you can thereby know this proposition. That appears to hold even when you simply have no reason to doubt my credibility.[11]

These points suggest a *principle of testimony-based justification*: At least normally, a belief based on testimony is thereby justified (i.e., justified on the basis of the testimony) provided the believer has adequate (situational) justification for taking the attester to be credible regarding the proposition in question. There is no easy way to specify the conditions for adequacy here. What we may say is that in everyday life people who do not often find people speaking falsely to them will have adequate justification for taking testimony they receive to be credible in the absence of special reason to doubt it.

We might formulate a similar principle for knowledge. To see its content let us speak of *undefeated testimony* when testimony occurs in the absence of at least the following common and probably most characteristic defeaters, i.e., factors that preclude testimony's giving the recipient knowledge: (1) internal inconsistency in what is affirmed, as when an attester gives conflicting dates for an event; (2) confused formulation, a kind that will puzzle the recipient and tend to produce doubt about whether the attester is rightly interpreted or even has a definite belief to communicate; (3) the appearance of insincerity, as when the attester seems to be lying; (4) conflict with apparent facts evident in the situation in which the testimony is given, as when a person shoveling earth over smoking coals says there has been no campfire; and (5) conflict with what the recipient knows, justifiedly believes, or is justified in believing (has justification for believing). These conditions may occur separately or together; the more of them an attestation satisfies, the more clearly defeated it is, other things being equal. Trust in testimony, I take it, frees us from drawing justificatory inferences whenever we receive testimony. But if we do not have justified trust—or at least justification for having trust—then we should not believe the attester and will not be justified in believing *p* if we do.

In this light, it is plausible to hold the *principle of testimony-based knowledge*: A belief based on undefeated testimony normally constitutes knowledge provided that the attester knows the proposition in question and the believer has no reason to doubt either this proposition or the attester's credibility regarding it.[12] Neither this principle nor the justification principle is unqualified, but there may be only a very few cases in which abnormal conditions prevent testimony from yielding justification or knowledge (or both) when the specified conditions are met.

The twofold epistemic dependence of testimony

Whatever we say about the exact conditions under which testimony grounds knowledge or justification in its recipient, we have so far found no reason to doubt that under *many* conditions testimony is a source of both knowledge and justified belief on the part of its recipient. It has seemed so far, however, that testimony cannot be a basic source of knowledge, since one cannot know something on the basis of testimony unless the attester knows it. This is why testimony does not, as such, generate knowledge though it may be described as transmitting it.

Testimony may, of course, generate knowledge *incidentally*, as when, by attesting in a surprised tone that it is 4 a.m., I give a fellow insomniac knowledge that I am awake. This knowledge is grounded not on my testimony but on the mere *hearing* of it. That kind of knowledge could as easily have been conveyed without testimony, by rising from my chair.

Testimony, like inference, can exist in indefinitely long chains. An attester might know that *p* on the basis of a third person's testimony, and the third might know it on the basis of testimony by a fourth, rather than from a generative source such as perception. But how far can this go, with each attester

informed by a previous one? There is surely some limit or other in each situation, as opposed to an infinite regress (difficulties with infinite regresses will be pursued in Chapter 9).

That brings us to a second respect in which testimony cannot be a basic source of knowledge. Surely if no one knew anything in a non-testimonial mode, no one would know anything on the basis of testimony either. More specifically, testimony-based knowledge seems ultimately to depend on knowledge grounded in one of the other sources we have considered: perception, memory, consciousness, and reason. To enable others to know something by attesting to it, I must know it myself; and my knowledge must *ultimately* depend at least in part on someone's non-testimony-based knowledge, such as knowledge grounded in seeing that the clock says four.

One might try to reinforce this view as follows. Even if someone had previously attested to a proposition, I would have to *perceive* this and to know some supporting proposition, say, that someone had credibly said it is four. But this claim is mistaken. The required kind of perceiving does not entail forming a belief of this sort, perhaps not even the specific perceptual belief that someone said it is four. The case shows, then, only that testimony is *operationally dependent* on perception, not that it is *inferentially dependent* on perceptual *belief*. It requires perceptual raw materials, but not beliefs of premises about those materials.[13]

If, as seems to be the case, testimony-based knowledge and justification do not depend on premises that support the testimony-based belief—say, premises confirming the credibility of the attester—this explains how such a belief can be basic. Testimony as a source of knowledge and justification need not be basic relative to other sources of knowledge and justification in order for beliefs grounded in it to be basic in the order of beliefs.

The point that testimony-based beliefs can be basic is entirely consistent with the point (made earlier) that the *attester's* knowledge that is the ground of the hearer's (potentially basic) knowledge cannot ultimately be based on testimony. Knowledge that is directly and wholly based on testimony for the recipient cannot be *ultimately* based wholly on testimony for the giver. The first would have no "right" to transmit it to the second, just as I would have no right to give someone what I had merely borrowed from someone else, who had merely borrowed it from a third person, and so on to infinity.

The point that testimony-based beliefs can be non-inferential and in that way not dependent on premises is important. But the operational dependence of testimony has both epistemological and conceptual significance. For if one did not have perceptual *grounds* for knowledge, or at least for justified belief, that someone has attested to the proposition in question, one could not know it on the basis of the testimony. This is an epistemic dependence not paralleled in the case of perception.[14] It shows that even if testimony-based knowledge need not inferentially depend on having *knowledge* grounded in another mode, it does epistemically depend on having *grounds*, from another mode,

grounds *for* knowledge in that other mode. Testimony-based knowledge thus depends on—and in this sense presupposes—the availability, or one might say the potential cooperation, of another source of knowledge, even if such knowledge does not require the actual operation of that source in yielding beliefs of the premises it stands ready to supply.

The case with justification is similar on this point. I cannot acquire justification for believing something on the basis of testimony unless I have justification for believing that the testifier is credible, as well as for certain other propositions, such as that I heard the testimony correctly. This justification cannot come entirely from testimony. Suppose Jane assures me about Bert, but I am not justified in taking Jane to be credible. Juan now tells me that Jane is utterly reliable. But how can this help unless I am justified in trusting Juan? Non-testimonial grounds of justification, such as perception of Juan's conduct or a memory of his guiding me in the past, must at least tacitly play some role in justifying my believing him. This role need not, however, be inferential: they need not produce in me beliefs of premises from which I infer that he is credible; they simply give me a justification that I can appeal to in framing such premises if I need them.

It may help to describe one of my overall conclusions—that testimony is not a basic source of knowledge or justification—as reflecting a disparity between the superficially simple psychology of testimony and its even more complex epistemology. Often, when we hear people attesting to various things, we just believe these things, non-inferentially and even unreservedly. But this natural psychological process yields knowledge and justification in the recipient only when certain epistemic conditions are met. In the case of testimony-based knowledge, there must be knowledge, even if not necessarily justification, on the part of the attester, whereas in the case of testimony-based justification there must be justification, even if not knowledge, on the part of the recipient. The first requirement concerns the attester's epistemic situation with respect to the proposition attested to; the second concerns the recipient's epistemic situation with respect to the attester, or the proposition, or both.[15]

The indispensability of testimonial grounds

The epistemic dependence of testimony on other sources of belief must be squared with the fact that tiny children learn from what others tell them even before they are properly said to be justified (or unjustified) in believing what they do. Consider teaching color words. After a time, the child learns that the sofa, say, is red. But the tiny child has no concept of credibility or other notions important in gaining justification from testimony and, initially, insufficient experience to be justified in believing that its adult teachers are credible. On the view developed here, however, this point is quite compatible with the child's acquiring certain kinds of *knowledge*.

Conceptual versus propositional learning

The first thing to note in explaining this compatibility is that there are at least two ways to learn from testimony: one can learn (in the sense of coming to know) the content attested to, and one can learn something *shown*, but not *stated*, by the testimony itself. The first case is learning *that*, specifically, that something is so. The second is learning *of* or *about* something (and may extend to learning how to do something). A tiny child just learning the basic colors is not, primarily, learning *that* (say) the sofa is red, but, above all, becoming aware of redness as the color of the sofa. This is learning colors and may be learning at least something *about* them.

In introducing the word 'red', then, the parent is only incidentally attesting to the proposition that the sofa is red. The point is to pair the word with an instance of what it stands for, with the aim of teaching that word (or, say, what the color red is), and the child can learn the main lesson without conceptualizing the sofa as such at all (something required for propositionally believing that the sofa is red). The former case of attestation—the *propositional testimony*—commonly results in propositional knowledge; we would thus have *propositional learning*. The parental introduction of vocabulary by attestation—*demonstrative testimony*—commonly results in *conceptual learning*.

It is important to see that the success conditions for the introductory function of language apparently require that for the most part the attestations be at least approximately true. A child cannot learn 'red' unless, in teaching the child English, a goodly proportion of the samples to which 'red' is applied are in fact red.[16] This does not of course show that most testimony is true, but it does imply that *if* communication is occurring when testimony is given to children, then one may reasonably assume that both attester and recipient have at some point benefited from a background in which a substantial proportion of attestations of a certain sort were true. How else can children be plausibly thought to have learned the language in which the communication occurs? This in turn supports the reasonableness, in everyday communicative situations, of taking testimony to be normally credible.[17]

At the time concepts are initially grasped in childhood, it may not be necessary that (propositional) belief and knowledge are acquired in every case. Conditions sufficient for conceptual learning may not be *automatically* sufficient for propositional learning. Belief and knowledge are, however, normally acquired at the time that concepts are initially grasped, even if conditions for mere conceptual learning are not necessarily sufficient for propositional learning.[18] Testimony easily produces both together. But if it cannot produce the conceptual learning without propositional learning, it can produce the latter without the former. It can be concept-producing, belief-producing (where some of the beliefs constitute knowledge), or both. The former case seems to be the more primitive, and the conditions for its possibility should not be taken as sufficient for the possibility of the latter.

It is very difficult to say when a child begins to form beliefs, as opposed to mimicking its elders by uttering things that, in adults, would express beliefs. Let us suppose both that belief-formation comes very early in life and that many of the first beliefs formed are based on what adults tell the child is the case. The child's defenseless credulity is a precondition for learning. Must this pose a problem for the epistemology of testimony suggested here?

Testimony as a primeval source of knowledge and justification

Very early in their lives we speak of babies and children as knowing things. One might object that this kind of talk is simply projective: *we* would know in their situation if we behaved in the relevant way, so why not say the child does? This is a defensible response, but suppose that at least by the time children begin to talk they do know certain things. We may surely speak of their *learning*—that the milk spills when tipped, that the stove is hot, and so on—and learning (in general) implies knowledge. At about the same time, children begin learning on the basis of testimony, say that steaming tea is hot and that when the doorbell rings, someone is outside.

If, as seems a reasonable assumption, gaining testimony-based knowledge requires only having no reason to doubt the credibility of the knowing attester, then the view proposed above encounters no difficulty. If a tiny child perhaps *can* have no reason for doubt, at least the child has none; nor need there *be* any reason, since much testimony is both undefeated and unassailably credible.

Suppose, however, that a stronger requirement must be met: that the child must have (possibly in a preconceptual way) some ground for taking the speaker to be credible, for instance a series of experiences repeatedly corresponding to what the speaker says. Perhaps we could sketch a case of having such a correlational ground that would be elementary enough to fit the rudimentary character of the child's knowledge. I doubt, however, that such a ground is required for testimony-based knowledge.

With justification, there may be greater difficulty in accounting for the case of tiny children. But the first thing to notice is that we do not use the vocabulary of justification, as compared with that of knowledge, for conceptually undeveloped creatures. For a child to be justified in believing that the sofa is red, the child would have to be capable not only of having a ground for believing this but also, correspondingly, of failing to have one and yet believing this proposition anyway, thereby being *un*justified.

It is arguable that by the time we may properly speak of children in this two-sided way as justified and also as unjustified—which is sometimes not long after they can speak—they *do* have a sense of the track record of adults in giving them information that their experience confirms. They have learned that if parents say it is cold outside, it is; and so forth. Children do not *use* the notion of credibility; but they can comprehend related concepts, such as

those needed for understanding that Mommy is right about things and baby brother must be corrected.

With testimony-based knowledge, by contrast, not even such unselfconscious justification seems required. The conditions by which knowledge is testimonially transmitted seem not to depend on justification in the recipient in the same way as does testimony-based justification. To be sure, testimony may be epistemically defeated—prevented from giving the recipient knowledge—by justified beliefs of some proposition contrary to the one attested to. But in the absence of defeaters, the recipient acquires testimony-based knowledge.

The acquisition of testimony-based justification seems to come later than that of testimony-based knowledge. One possibility for explaining how, very early in life, children may acquire an elemental kind of justification for accepting testimony is that at a very early stage they acquire a sense that they *themselves* generally give information only when they have gotten it themselves, as when they see that it is snowing or they feel hungry. For misinformation, we commonly and sometimes sternly correct children, whereas we patiently instill habits of correct reporting. This correlational sense that children apparently develop might provide a kind of analogical justification for taking others to be providing, when they give testimony, information *they* have obtained. A related and complementary hypothesis is that children have a rudimentary understanding of others in terms of what apparently explains their observed behavior. And what would explain Mommy's saying that it is snowing outside as well as her having seen that it is?

None of this is to say just when knowledge or justification enters the scene in human development, whether through the basic sources or through testimony. These are psychological questions; a philosophical theory need only leave room for plausible answers to them. The theory given here suggests that knowledge may arise before justification, but it does not entail even that. Moreover, it has at least this much harmony with the most familiar data about human development: the more natural it is, and the less figurative it seems, to speak of growing children as acquiring knowledge and justification based on testimony, the easier it is to find some elementary way in which they can satisfy the epistemic and justificational conditions set out above, such as making discriminations that enable them to assess what they are told and gaining some sense of the track record of those around them who offer information.

To say that testimony is not a basic source of justification or knowledge is not to imply that it is any less important in normal human life than a basic source. A source of knowledge and justification can be indispensable in life even if it is not basic. It may be that no normal human being would know anything apart from dependence on receiving testimony.[19] If there is no innate knowledge, and if we know nothing before learning a language (something I here assume for the sake of argument but wish to leave open), then unless we could acquire linguistic competence without the help of others, they would be essential in our coming to know anything at all.

If we try to imagine what would be left if we gave up all the knowledge and beliefs we have acquired on the basis of testimony, we would be quite unable to accomplish the sorting in the first place. But even beginning the task of trying to put aside what one knows on the basis of testimony suggests that one would at best be thrust back to a primitive stage of learning. I want to pursue this idea in relation to David Hume.

Non-testimonial support for testimony-based beliefs

If one ponders Hume's view of testimony as capable of grounding knowledge only on the basis of a kind of validation by other sources, one may want to know to what extent testimony-based knowledge and justification, even taken item by item, can be backed up by other kinds. For Hume, our "assurance" in any matter depending on testimony "is derived from no other principle than our observation of the veracity of human testimony, and of the usual conformity of facts to the reports of witnesses."[20]

Let us ask whether, for each proposition one justifiedly believes (wholly) on the basis of testimony, one has a justification from other sources. Call this the *focal justification question* for testimony-based beliefs. We must immediately acknowledge a complicating factor. These other sources would include propositions one justifiedly believes on the basis of memory; and although one's justification for these propositions would not depend on the testimony needing support, one's beliefs of them might have been *originally* based on testimony. Much of what is stored in memory we came to believe through what others have told us in person or in writing. Still, we might be memorially justified in holding the beliefs in question even if we have forgotten their testimonial origin. Suppose, however, that we reasonably believe—as we should—that many of these beliefs have arisen through testimony. Then if we do not generally trust testimony, that reasonable belief might reduce our justification for the beliefs in question. For all that, if what we are memorially justified in believing, together with other justified beliefs we hold on the basis of our non-testimony-based experiences, can justify believing numerous propositions we find affirmed in testimony, then perhaps we do have *some* independently grounded justification for everything we justifiedly believe on the basis of testimony.

Given that memory is a basic source of justification, then, we might have memorial justification for beliefs which only *seem* to be grounded in actual past experience. In any case, many of our beliefs about conditions under which people are credible are preserved in, or at least justified by, our memories. Thus, even if I have no evidence regarding *p* I may have reason to think the attester's affirming it is some reason to believe it.

To be sure, some of the memorially justified beliefs in question would not be justified unless I had been at some point justified in believing something on the basis of testimony, as when I believe one person's testimony, remember the proposition attested to, and use it in checking on another person's testimony. There may be a kind of circularity here, since testimony plays a

role in checking on the credibility of testimony. But notice two points. First, there are *two* attestations, normally by *different* attesters. Second, it might be argued that since memory is a basic source of justification, it may yield justification that supports testimony but is not testimony-*based*. Even if a memorially justified belief is originally justified on the basis of testimony, it may later be justified without dependence on that initial justification.

To illustrate some of these points about justifying a belief based on testimony, take a case regarding a country I do not know first-hand. Consider a radio news program announcing an earthquake in Indonesia. I have—though I may never have articulated it—a sense of the track record of the network in question and of the geology of Indonesia, a sense of how often errors of that kind are made, and so forth. One could always say that this yields a very weak justification, especially since I rely on some beliefs acquired through testimony (though that testimony may be independent of the credibility of the network in question). Certainly such a justification is far from conclusive. But there is still no good reason to think it must be inadequate.

It is natural here to raise a further question, a *global justification question*: Could one fashion an overall justification of the entire set of the propositions one believes on the basis of testimony? There are at least two questions one could be asking here. If the reference is to all the propositions one believes conjoined together—to the "long" proposition consisting of the first *and* the second *and* the third item, etc.—then one cannot even imagine contemplating such a monstrosity, much less justifying it. If, however, the reference is to the *set* of one's testimony-based beliefs considered in the abstract, it is still not clear how to conceive justifying it. Suppose we take it to be a matter of showing that "by and large" testimony-based beliefs are justified. If we do not allow some testimony-based beliefs to justify others and we try to suspend judgment on all such beliefs we hold (assuming such massive suspension of judgment is even possible), I do not see that this corporate global justification project would work.[21] Let me explain.

Whatever might be possible in principle, it is doubtful that we can always avoid relying on testimony, at least indirectly, in any actual appraisal of testimony. Even one's sense of an attester's track record, for instance, often depends on what one believes on the basis of testimony. Think of how one news source serves as a check on another: in each case, testimony from one source is tentatively assumed and checked against testimony from another. How, then, can we globally justify testimony if we can never rely on it in the process?

There seems, moreover, not to be any general procedure by which one can produce a global justification for the proposition that the whole set of one's testimony-based beliefs (or even a major proportion of its elements) is justified. Fortunately, that project of global justification is not one we need attempt, and the epistemology of testimony I have sketched implies on this matter at most that justified testimony-based beliefs are, to some degree, *individually* justifiable for the believer in terms of the basic sources of belief.[22]

Some are thus justifiable, even if it turns out that not all of them are.

Sometimes one person can confirm another's testimony simply by observing the scene described in the testimony. Indeed, the reliability of testimony, whether on a particular occasion or in a general way, can be checked through the basic sources. This is significant, for it appears that no parallel point holds for the basic sources. For instance, I cannot check on the reliability of perception, either in a particular instance or, especially, in a general way, without appealing to that very source, as when I look at something again in better light to check my color judgment. (One can use data from one sensory mode to justify a belief arrived at in another, but this is still relying on one perception to check on the reliability of perception.) Nor can I check on the reliability of memory, say by revisiting the scenes of past experiences, without presupposing that I remember the original judgments I seek to confirm by my visitation. Similar points hold for self-consciousness and (intuitive) reason. This contrast is one reason testimony is not fully on a par with the basic sources: we can confirm or disconfirm testimony without relying on it, whereas we cannot confirm a deliverance of a basic source without relying on (at least) that very source. The contrast does nothing, however, to suggest that in human life as we know it, testimony is not essential for at least a huge amount of what we know.

Testimony is a pervasive and natural source of beliefs. Many testimony-based beliefs are justified or constitute knowledge. They may even constitute basic knowledge or basic belief, both in the sense that they are not grounded in premises and in the sense that they play a pivotal role in the life of the believer. We might thus say that testimony-based beliefs not only constitute some of our basic knowledge but also are psychologically and existentially basic.

These beliefs are, however, not unqualifiedly basic epistemically. They are basic only in the sense that they are not *inferentially* dependent on knowledge or justified belief of prior premises. They are *epistemically* dependent, in a way perceptual beliefs are not, on one's having grounds for knowledge or justification, and they are psychologically dependent on one's having some ground—such as hearing someone speak—in another, non-testimonial experiential mode. Testimony-based beliefs, then, are not premise-dependent but do depend, for their epistemic or justificational status, on the basic experiential sources of knowledge and justification considered in Chapters 1–6. As a source of knowledge and justification, testimony depends both epistemically and psychologically on these other sources. This is entirely consistent, however, with its playing an incalculably important role in the normal development of our justification and knowledge.

Notes

1 For a wide-ranging, historically informative account of what constitutes testimony for and numerous epistemological problems surrounding it see C.A.J. Coady, *Testimony* (Oxford: The Clarendon Press, 1992).
2 Not all testimony is *verbal*, much less *oral*. Consider someone's asking

of a person who requested testimony on a crime, 'What did he say?' A perfectly good answer would be 'That Mack the Knife did it', even if this was affirmed by sadly nodding when asked whether Mack is the one who did the deed. The concept of testimony allows numerous ways of telling people things; certainly any symbolic behavior rich enough to count as affirming a proposition can serve.

3 Thomas Reid spoke eloquently on this topic; he said, for example, "The wise author of nature hath implanted in the human mind a propensity to rely upon human testimony before we can give a reason for doing so. This, indeed, puts our judgment almost entirely in the hands of those who are about us in the first period of life." See 'Essay on the Intellectual Powers of Man' in *Thomas Reid's Inquiry and Essays*, ed. by Ronald Beanblossom and Keith Lehrer (Indianapolis: Hackett, 1983), p. 281.

4 Granted, I must have (and so must memorially retain) a concept of a tree; but this merely conceptual memorial state is not a potential source of belief (which is not to deny that it can play any other kind of causal role in belief-formation).

5 Three points may help here. First, telepathic or otherwise strange reception of testimony may, at least for our purposes, be construed as some kind of perception. Second, granting that one cannot form perceptual beliefs without having whatever additional beliefs may be needed to possess the concepts required to understand what is believed perceptually—in my example, for instance, the concept of a star-gazing stroller—this does not imply the kind of dependence on any other belief source exhibited by that of testimony upon perception. One can perceive, though not interpret, such a stroller without having these concepts; one cannot even receive testimony, and so cannot begin to interpret or learn from it, without perceiving it. Third, supposing perception cannot occur without some manifestation in consciousness (which is itself a source of beliefs), here consciousness is an element in perception in a way that perception by an audience is plainly not an element in testimony.

6 You might come to know it on the basis of something *about* my testimony: perhaps, for example, I give it nervously and you know that the nervousness is an after-effect of my being shaken by the fit of temper which I have since half forgotten and might even deny. This would be a case of belief *caused* by testimony but not *based* on it (not an easy distinction to explicate; but it was illustrated in Chapter 3 in showing how a belief that a past event occurred need not be a memory belief even if caused by that event; and it will be developed further in Chapter 10). One requirement for a belief to be based on testimony is the believer's holding the proposition because it was attested to, as opposed, for example, to holding it because of *how* or from what motive it was attested to. There have, however, recently been challenges to the idea that testimony-based belief constitutes knowledge only if the attester knows the proposition in

question. Some of these are cited and briefly answered in my 'Testimony, Credulity, and Veracity', in J. Lackey and E. Sosa (eds.), *The Epistemology of Testimony* (Oxford: Oxford University Press, 2006).

7 The qualifier 'testimony-based' is crucial: suppose I attest, in a baritone voice, that I have a baritone voice, but do not know this because I falsely believe I have a tenor voice; then you come to know, *from* my testimony, but not on the basis of it (its content), that the proposition to which I attest is true. The same point holds for justification in place of knowledge. One might also say that you come to know it *through* my testimony in a weak sense of 'through' not implying that the content of what I attest is crucial. It is also possible that the content, but not *my attesting it*, is essential, as when I present an argument you know I barely understand, and you come to know its conclusion, not because I attest to it or to the premises, but on the basis of yourself realizing, by bringing to bear your background knowledge, that they are true and entail the conclusion. This would be knowledge based on the *content* of testimony, but it would not be what we call 'testimony-based knowledge'.

8 I develop and defend this contrast in 'Memorial Justification', *Philosophical Topics* 23 (1995), 251–72. Particularly interesting from the point of view of the thesis that the attester must know that *p* are examples given by Peter Graham, 'Conveying Information', *Synthese* 123 (2000), 365–92 and Jennifer Lackey, 'Testimonial Knowledge and Transmission', *Philosophical Quarterly* 49 (1999), 471–90. I will mention just one of hers. A teacher who disbelieves the theory of evolution but teaches it conscientiously tells his students, on the basis of his correct reading of the theory and his observation of a fossil, that there were *Homo sapiens* in the place in question. Since we may suppose he is giving his students correct information on a sound basis, we may tend to conclude that testimony-based belief (theirs) can be knowledge without the attester's knowing the proposition in question. This is an interesting case, since the hearers do have a testimony-based true belief that seems well grounded. But is it knowledge, if the teacher would have taught a false theory in the same way, had his job required it? Even if the theory itself is (an item of) "knowledge" (as some would say if it is known), *he* isn't a reliable link in the chain from the fossil through the theory, since he neither knows it nor even believes it on the kind of ground that would protect him from error in the way the (truth-conducive) grounds of knowledge do. It isn't that the theory he uses just happens to be true, but—from the point of view of genuine evidence—he just happens to use it. If, on the other hand, we suppose that the school would not require teaching a theory that is not well evidenced, and that the students believe something to this effect, then perhaps an essential part of their basis for believing him is that background belief. Their belief would then be bolstered by background beliefs rather than a genuinely testimony-based

one. It would be as if they had to believe something to the effect that this is what the school is teaching in order to believe what he says. Chapter 10 will discuss knowledge in a way that supports this analysis.

9 I leave open whether knowledge transmitted by testimony can be as *well grounded* as that of the attester (though I am inclined to think it can be, say when the attester is "absolutely" reliable, a property that in principle could perhaps belong to memory in some cases). By contrast, so far as knowledge goes, "a testimonial chain is no stronger than its weakest link," as Alvin Plantinga puts it in *Warrant and Proper Function* (Oxford: Oxford University Press, 1993), p. 84. He is speaking of what he calls *warrant*, which he views as roughly what makes true belief knowledge; and if, as I suspect, the point holds there too, then justification contrasts with warrant on this score as it does with knowledge. It should be added that if knowledge cannot be stronger than its weakest link, it probably need not be any weaker.

10 If this is so, it may show something else: on the assumption that you cannot know a proposition on the basis of premises you do not also know, this case would show that your testimony-based knowledge is not inferential, since the would-be credibility premise is not known but only permissibly assumed.

11 One possibility raised here is that of knowledge without justification. This will be considered in some detail in Chapter 10.

12 These principles are formulated cautiously in several ways: for instance, they allow for abnormal circumstances to provide exceptions; they allow that the resulting justification not be strong but only "adequate" for reasonable belief; they allow, but do not entail (what I think plausible but leave open), that the testimony-based belief *always* acquires prima facie justification from the testimony; and they permit the recipient to have justification or knowledge of the proposition in question from some *other* source as well. The epistemic principle might well be broadened by specifying that the recipient has no *overall* reason for doubt, but I offer that as a suggestion without adopting it.

13 Here I differ from Elizabeth Fricker, who (in one place) maintained that the recipient must perceptually believe "that the speaker has made an assertion with a particular content . . . capable of being knowledge . . . I have been convinced by John McDowell's contention that hearers' perceptions of speakers' utterances are . . . a case of perceptual knowledge." See 'The Epistemology of Testimony', *Proceedings of the Aristotelian Society* 61 (1987), 70. The reference to McDowell is to 'Anti-realism and the Epistemology of Understanding', in H. Parret and J. Bouveresse (eds.), *Philosophical Subjects* (Oxford: The Clarendon Press, 1980).

14 John Greco (in correspondence) has raised the question why consciousness is not related to perception as perception is to testimony. The beginning of an answer may be that (1) (sensory) consciousness of a ground for *p* is a *constituent* in perception that *p*, whereas no ground for *p* is a

constituent in testimony that *p*; (2) perceptual justification and knowledge entail and depend on *having* consciousness of the perceptual object which the justification or knowledge concerns; whereas (3) testimony-based justification and knowledge do not entail or depend on perception regarding what the justification or knowledge concern—namely, the proposition attested to (or even its subject matter, with which the recipient may have no relevant experience). Testimony-based justification, moreover (though not testimony-based knowledge), also normally depends on perception (or at least on sensory experience) *separate* from that required to receive the testimony; for (on the view taken in this chapter) justification for accepting the credibility of the testimony normally requires perception (or at least sensory experience) as part of the background one needs to acquire testimony-based justification. (Note 19 indicates why the normality qualification is needed here.)

15 The epistemology of testimony suggested here may perhaps be more stringent than that of Thomas Reid. For an interpretation and defense of the apparently Reidian view that testimony-based beliefs need not depend even for their justification on other sources of justification see Mark Owen Webb, 'Why I Know about as Much as You: A Reply to Hardwig', *Journal of Philosophy* 90 (1993), 260–70.

16 Strictly, the samples need only look red, as when white objects are flooded by red light; and arguably, one could even teach 'red' by producing only hallucinations of the color.

17 It can be connected with arguments such as we find in Donald Davidson's work for the conclusion that most of our beliefs must be true, but it does not imply that stronger conclusion. For discussion of this and other Davidsonian hypotheses, see Coady, *Testimony*, esp. Chapter 9. Cf. Tyler Burge's view that "We are a priori prima facie entitled to accept something that is prima facie intelligible and presented as true." See his 'Content Preservation', *Philosophical Review* 102 (1993), 472. Some explication and discussion of this view is provided in my 'Testimony, Credulity, and Veracity', cited above.

18 It is difficult to see how one could, through testimony, produce conceptual learning without producing some belief. Could a child become acquainted with what redness is in connection with being told the sofa is red, yet not acquire a belief of some sort, for example objectually believing the sofa to be red? There is no need to settle this matter here; nor can I pursue related questions concerning conceptualization in higher animals.

19 One reason this point is restricted to normal human beings is that it seems possible for a human being to be created, as a full-blown adult, artificially, in which case much knowledge of abstract propositions and perhaps of other sorts, such as knowledge of the perceptible external environment in which the person is made, can occur before any testimony enters the picture. The story of Adam and Eve is a theological

version of creation at the adult stage. There are also evolutionary conceptions of how knowledge first arises in human history, but these genetic questions would take us too far from our main questions.

20 *An Enquiry Concerning Human Understanding*, ed. by L.A. Selby-Bigge (Oxford: Oxford University Press, 1888), section 88.

21 We would certainly not be able to appeal to any significant segment of scientific knowledge, for there we are heavily dependent on testimony, written and oral. A plausible case that this dependence is even greater than it seems is made by John Hardwig in 'Epistemic Dependence', *Journal of Philosophy* 82 (1985), 693–708.

22 For supporting considerations favoring the possibility of the local justification and opposing that of a global one, see Elizabeth Fricker, 'Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony: C.A.J. Coady's *Testimony: A Philosophical Study*', *Mind* 104 (1995), 393–411.