

Cambridge University Press
978-0-521-62496-1 - Decision Theory as Philosophy
Mark Kaplan
Frontmatter
[More information](#)

It has been over sixty years since Ramsey first argued for the significance of decision theory to epistemology. Yet many philosophers remain unconvinced. The familiar probabilistic constraints decision theory imposes on opinion and confirmation seem too demanding, too prudential in nature and too tangential to our concern with the propriety of categorical belief and knowledge.

How important are these concerns? Mark Kaplan argues that they are very important – but that a properly modified and interpreted Bayesian decision theory can meet them. His brief is that, suitably formulated, Bayesian decision theory is of the most profound philosophical consequence to the way we are accustomed to think about inquiry, criticism and rational belief.

Kaplan makes his case in a clear and compelling way, and with a minimum of technical detail. The modest variant of Bayesian decision theory to which he appeals is new, well-motivated and easy to follow. Most proofs are relegated to an appendix. A brief primer on probability is also provided. This book is not just an original contribution to Bayesian epistemology. It is also the most accessible treatment available of the relation Bayesian epistemology bears to the rest of the field.

Cambridge University Press
978-0-521-62496-1 - Decision Theory as Philosophy
Mark Kaplan
Frontmatter
[More information](#)

Decision Theory as Philosophy

Cambridge University Press
978-0-521-62496-1 - Decision Theory as Philosophy
Mark Kaplan
Frontmatter
[More information](#)

Decision Theory as Philosophy

Mark Kaplan
University of Wisconsin–Milwaukee



Cambridge University Press
 978-0-521-62496-1 - Decision Theory as Philosophy
 Mark Kaplan
 Frontmatter
[More information](#)

Published by the Press Syndicate of the University of Cambridge
 The Pitt Building, Trumpington Street, Cambridge CB2 1RP
 40 West 20th Street, New York, NY 10011-4211, USA
 10 Stamford Road, Oakleigh, Melbourne 3166, Australia

© Cambridge University Press 1996

First published 1996
 First paperback edition 1998

Printed in the United States of America

Library of Congress Cataloging-in-Publication Data

Kaplan, Mark, 1951–
 Decision theory as philosophy / Mark Kaplan.
 p. cm.

Includes bibliographical references.

ISBN 0-521-47505-8

1. Decision-making. 2. Knowledge, Theory of. 3. Science –
 Philosophy. I. Title.
 BD184.K37 1995 96-868
 121'.6 – dc20 CIP

A catalog record for this book is available from the British Library.

ISBN 0-521-47505-8 hardback
 ISBN 0-521-62496-7 paperback

Cambridge University Press
978-0-521-62496-1 - Decision Theory as Philosophy
Mark Kaplan
Frontmatter
[More information](#)

To my parents, Edward and Sue Kaplan

Contents

<i>Preface</i>	<i>page ix</i>
<i>Note to the Reader</i>	xv
CHAPTER ONE: Confidence	1
I. A Decision Problem	1
II. The Five Principles	4
III. Decision Problem Solved	14
IV. Modest Probabilism	20
V. The Sin of False Precision	23
VI. Countable Additivity	32
VII. Philosophical Objections	36
CHAPTER TWO: Evidence	45
I. A Puzzle about Evidence	45
II. Evidence and Confidence	47
III. Puzzle Solved	52
IV. “Popperian Animadversions”	56
V. Evidence and Conditionalization	62
VI. An Argument for Conditionalization	67
VII. The Argument Disarmed	70
VIII. Reflection	73
IX. Evidence: Quality vs. Quantity	75
X. Subjectivism	85

CONTENTS

CHAPTER THREE: The Bayesian Challenge	89
I. What about Belief?	89
II. The Certainty View	91
III. The Confidence Threshold View	93
IV. The Bayesian Challenge	98
CHAPTER FOUR: Rational Belief	102
I. The Bayesian Challenge Revisited	102
II. Acting on Beliefs	103
III. The Assertion View	107
IV. Deductive Cogency: A Puzzle	112
V. Structural Soundness	121
VI. Adjudicating between Theories	132
VII. Global Interests	136
VIII. Puzzle Solved	140
IX. Believing the Improbable	142
X. The Assertion View: Objections	148
CHAPTER FIVE: The Bayesian Canon	155
I. The Dutch Book Argument	155
II. Ramsey's Utility Theory	161
III. Degrees of Belief Constrained	168
IV. Three Objections	172
V. Savage and Jeffrey	176
CHAPTER SIX: Decision Theory as Epistemology	181
I. Staking the Claim	181
II. The Naturalists Object	183
III. A Methodological Moral	186
APPENDIX 1: Principles and Definitions	191
APPENDIX 2: Proofs	199
APPENDIX 3: Probabilism – Some Elementary Theorems	212
<i>Bibliography</i>	215
<i>Index</i>	223

Preface

It has been a traditional concern of epistemology to find what Descartes called “rules for the direction of the mind”: a set of principles that will say, in some general yet useful (although, perhaps, not exhaustive) way, how the opinions of a rational person ought to be constrained. Epistemologists have sought such principles, and evidence for their legitimacy, in various places. Some have sought insight from the traditional sources – from reflections on the metaphysical structure of the world, from analyses of the nature of justification. Others have argued that insight is available only from the scientific study of the empirical world – from the analysis of the way actual human inquirers behave, from the findings of cognitive psychology, from the application of evolutionary biology to human cognition, from the study of artificial intelligence.

What sets the Bayesian approach to epistemology apart from the rest is that its proponents look in a different place. They look for rules for the direction of the mind in the theory of rational preference – in decision theory. At first blush, this looks like a preposterous undertaking. After all, one would think that (if anything) it is epistemology that would place constraints on what we are rational to prefer, not the other way round. But there is method in the Bayesian madness.

Suppose you are offered a free choice between two identical gambles, one a gamble on the hypothesis h and the other a gamble on the hypothesis g . The prizes in the two gambles are so designed that you would be just as happy winning the one gamble as the other, just as unhappy to lose the one as the other. You would, of course, prefer to win. Now, think about what will decide the choice for you. It is pretty clear, isn't it, that your preference will depend entirely on

PREFACE

which (if either) of the two hypotheses you hold to be more probably true: you will find the gamble on b preferable to the gamble on g if and only if you are more confident in the truth of b than you are in the truth of g .

This insight is modest to be sure. But its significance is not. For this insight suggests that any theory that says something general and informative about when it is rational for you to prefer one gamble to another stands some chance of saying something about when it is rational to invest more confidence in one hypothesis than in another. The insight suggests that, should we be able to come up with a credible decision theory, we may well find that we have come up with a theory from which we can learn something interesting about the way in which we should invest confidence in hypotheses.

But *is* there a credible decision theory capable of teaching us anything of genuine epistemological interest? The answer is a matter of some controversy. Ever since the 1920s, when Bayesian decision theory received its first sophisticated expression in the work of Frank Ramsey, there have been those who have argued that the answer is “Yes.” And for almost as long, others have maintained that neither Ramsey’s efforts nor the efforts of those who have followed him are of any genuine epistemological consequence.

I have some sympathy for the nay-sayers. For example, I find it wildly implausible to suppose (as some orthodox Bayesians have done) that actual investigators in general harbor precise, real-valued degrees of confidence for hypotheses. Even when construed as a regulative ideal, the requirement that investigators harbor such precise degrees of confidence looks as if it owes more to an unfortunate worship of false precision than it does to reason. I also see a great deal of strain in some of the arguments with which Bayesians have tried to convince us that rational degrees of confidence are subject to probabilistic constraint. It seems to me that the Dutch Book Argument, the one most often produced when there are philosophers in the audience, is patently unconvincing. Finally, I think that Bayesianism’s capacity for solving philosophical problems is much exaggerated by its proponents. In particular, the Bayesian claim to have provided a credible measure of the degree to which hypotheses are confirmed by evidence seems to me just false.

But for all this, I am convinced that those who would have us reject decision theory’s claim to philosophical importance are mistaken. My aim in this book is to say why. I mean to show how, from

PREFACE

the body of theory that Ramsey produced and inspired, one can extract an accessible and entirely compelling decision theory which (though modest in scope) has the most profound consequences for the way in which we are accustomed to think about, and conduct, the enterprise of inquiry, criticism and justification.

I recognize that many readers may find it hard to see how such a project could possibly succeed. Contemporary writers on inquiry, criticism and justification – epistemologists and philosophers of science – are worried about what we categorically *believe*. When do beliefs count as justified? Under what circumstances does a belief qualify as knowledge? How can a rational person believe a scientific theory true, if at all? Even if there is a compelling decision theory that avoids the excesses of orthodox Bayesianism, it would seem from what I have so far said that the most it could do is place constraints on how *confident* we are in the truth of hypotheses – surely a matter of marginal interest.

But it is precisely this comfortable assumption – that categorical belief is central to our pursuit of inquiry, criticism and justification – that decision theory calls into question. Indeed, I will be arguing that what makes decision theory such an important and disturbing contribution to epistemology is that it reveals just how unclear it is that there is *anything* the familiar questions and answers of epistemology and the philosophy of science (couched, as they are, in terms of categorical belief) can coherently be said to be *about*. Decision theory challenges the very intelligibility of the way we are most accustomed to think about our opinions.

Some Bayesians would go further. They would say that decision theory teaches us that we must abandon all talk of categorical belief and knowledge and focus our attention on states of confidence when we pursue our epistemological queries. But this, I think, is a mistake. There can be no question (as I hope to convince the readers of this book) that, both as we conduct inquiry and reflect philosophically upon it, we should be paying far more attention to our states of confidence. But, I will argue, there is also, salvageable from the wreck of our ordinary way of thinking about belief, a thin notion of categorical belief we cannot afford to ignore.

Indeed (I will argue) it is a notion of belief critical to Bayesianism's claim to epistemological import. Without it, Bayesianism cannot make sense of our interest in theory. With it, Bayesianism can do that and more. Bayesianism can provide an account of rational belief

PREFACE

that dissolves paradoxes (the preface and the lottery) and skeptical worries (such as those anti-realists have leveled against the rationality of believing in the truth of scientific theories) that have dogged every attempt to make sense of our commitment to theory.

I doubt, however, that anything I have so far said will have dispelled the skeptical worries harbored by readers who are sympathetic to the naturalist approach to epistemology. Central to that approach is the conviction that epistemology must take the form of a scientific inquiry that, to quote from a recent survey and defense of the approach (Kitcher 1992, pp. 75–6), “is to be carried out by describing processes that are reliable, in the sense that they would have a high frequency of generating epistemically virtuous states in human beings in our world.” This is a constraint that decision theory patently fails to satisfy. Far from a product of empirical research, Bayesian decision theory is the result of armchair reflection on the nature of rational preference. Thus, by naturalist lights, it is still hard to see how decision theory could constitute an important contribution to epistemology.

I will argue, however, that the naturalists are in error; false to their own practice of epistemic appraisal (and ours) and crippling to the enterprise of inquiry if adopted, the naturalists’ constraint on an adequate epistemology must be rejected. This is not to say that questions about the extent to which decision theory can impose constraints on actual human beings in our world are not worth asking. On the contrary, these are serious questions. But, as I will show, they admit of quite satisfactory answers.

I have so far only sketched the main argument that threads its way through the book. A lot has been left out. One chapter of the book is devoted to an extended discussion of the ways in which your states of confidence do (and do not) bear on how you appraise evidence, another to understanding how the decision theory I will be advancing is related to some of its famous Bayesian antecedents. There is also the matter (taken up at various points in the book) of showing why those Bayesian doctrines I have disparaged above are, indeed, mistaken.¹ But I think I have said enough to give the reader a fair

1. There is actually no place in the book at which I explicitly argue for the implausibility of the assumption that actual investigators harbor precise, real-valued degrees of confidence for hypotheses. But, for such an argument, see Kaplan 1989, pp. 48–55.

PREFACE

idea of what the book is about. I want now to acknowledge my debt to those who, like it or not, are in some way responsible for the way the book turned out.

In both style and substance, the book is much better for the benefit of helpful criticism from Brad Armendt, David Christensen, Gary Ebbs, Peter Haddawy, Richard Jeffrey, Hilary Kornblith, Julius Sensat, Jamie Tappenden, Bas Van Fraassen and William Wainwright, each of whom read, and offered comments on, one or more parts of the manuscript. It has also benefited from the suggestions offered, in conversation and correspondence, by Stewart Cohen, Ken Gemes and Mark Lance. But the book was positively transformed by Patrick Maher and James Joyce. Maher's trenchant criticisms of the first draft, and Joyce's comments (rich with both criticism and editorial insight) on a large portion of the penultimate draft, provoked me to make major revisions and improvements. To all of these individuals (and to the two careful and sympathetic readers for Cambridge University Press – one of whom revealed himself to be Peter Vallentyne) I extend my thanks.

The generous support of three institutions enabled me to finish the book when I did. The National Science Foundation funded my research for the year during which most of the first draft was written. The Philosophy Department at the University of Pittsburgh, which appointed me Visiting Scholar for the spring of 1994, provided office space, supplies and hospitality during the semester in which most of the final revisions to the manuscript were completed. The University of Wisconsin–Milwaukee awarded me the sabbatical leave that allowed me to spend that semester in Pittsburgh. I am grateful to all three for their help.

I come, finally, to a couple of special debts.

The first is to Richard Jeffrey and Isaac Levi, not so much for their acts of kindness toward me over the years (which have been many), but for the rich body of work they have made available to all. This book is very much a result of my attempt to come to terms with the vivid, profound and incompatible accounts of inquiry they have promulgated. In the pages below, the influence of their writings is pervasive.

The second is my debt to Joan Weiner, who read every draft, every bit, and every piece I wrote and was consulted on every problem of consequence that cropped up in their writing. The combination of incisive criticism and unflagging encouragement she provided from

PREFACE

the very first is as responsible as anything for this book's being written.

I have thought for many years about the matters with which this book is concerned. Some of those thoughts have already found their way into print. And, although I have changed my mind on some issues, fragments of the following essays are discernible in the pages that follow: "A Bayesian Theory of Rational Acceptance," *The Journal of Philosophy* 78 (1981): 305–30 (chapters 3 and 4); "Rational Acceptance," *Philosophical Studies* 40 (1981): 129–45 (chapters 3 and 4), © D. Reidel Publishing Company and reprinted by permission of Kluwer Academic Publishers; "Decision Theory as Philosophy," *Philosophy of Science* 50 (1983): 549–77 (chapter 1); "Bayesianism Without the Black Box," *Philosophy of Science* 56 (1989): 48–69 (chapter 1); "Confessions of a Modest Bayesian," in Jocelyne Couture and Kai Nielsen (eds.), *Reconstructing Philosophy? New Essays in Metaphilosophy, Canadian Journal of Philosophy Supplementary Volume* 19 (1993): 315–37 (preface, chapters 1 and 4); "Not by the Book," *Philosophical Topics* 21 (1993): 153–71 (chapters 1 and 5); "Epistemology Denatured," in Peter A. French, Theodore E. Uehling, Jr. and Howard K. Wettstein (eds.), *Midwest Studies in Philosophy XIX: Philosophical Naturalism* (Notre Dame: Notre Dame University Press, 1994): pp. 350–65 (chapter 6), © 1994 by the University of Notre Dame Press and used by permission; "Believing the Improbable," *Philosophical Studies* 77 (1995): 117–46 (chapters 3 and 4), © Kluwer Academic Publishers and reprinted by permission of Kluwer Academic Publishers. I thank the publishers of the journals and book in which these essays appeared for permission to reprint parts of them here.

For the paperback edition, I have corrected various typographical errors, expanded slightly my description of Case III on p. 30, and corrected an error in my characterization of the utility assignments displayed on p. 167.

Note to the Reader

My aim in writing this book was to reach a wide philosophical audience, including those who have no familiarity with decision theory or probability. Accordingly, the formal demands the book makes are minimal. They will be easily satisfied by anyone who has even a passing acquaintance with high-school algebra and truth-functional logic. Almost all the proofs are relegated to appendix 2, and the few elementary facts about probability that one needs to know can be gleaned by a quick perusal of appendix 3. It is also worth noting that those readers who are primarily interested in finding out what consequences decision theory may have for the way we think about justified belief and theory choice can satisfy their curiosity without reading the book cover to cover. They can skip chapter 5 and all but section II of chapter 2 without missing anything essential.