

THE ENTERPRISE OF KNOWLEDGE
An Essay on Knowledge, Credal Probability,
and Chance

Isaac Levi

The MIT Press
Cambridge, Massachusetts, and London, England

Publication of this volume has been aided by a grant from the
National Endowment for the Humanities.

To Judith

© 1980 by The Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any
form or by any means, electronic or mechanical, including photo-
copying, recording, or by any information storage and retrieval sys-
tem, without permission in writing from the publisher.

This book was set in VIP Times Roman
by DEKR Corporation
and printed and bound by Halliday Lithograph Corporation
in the United States of America.

Library of Congress Cataloging in Publication Data

Levi, Isaac, 1930-

The enterprise of knowledge.

Bibliography: p.

Includes index.

1. Knowledge, Theory of. 2. Probabilities. 3. Chance.
4. Inference (Logic). 5. Science—Philosophy. 6. Decision-
making. I. Title.

BD161.L378 121 80-10559

ISBN 0-262-12082-8

CONTENTS

	PREFACE AND ACKNOWLEDGMENTS	xiii
1	FALLIBILITY AND CORRIGIBILITY	1
	Knowledge without Pedigree	1
	Knowledge as a Standard for Serious Possibility	2
	Incorrigibility	5
	Truth	8
	Comparative Statics	9
	Infallibilism	13
	Truth as a Value in Inquiry	19
	Four Types of Revision	25
	True Justified Belief	28
	Free Speech	31
	Is Revision Possible?	33
2	TRUTH, INFORMATION, AND EXPANSION	34
	New, Error-Free Information	34
	Routine and Inferential Expansion	35
	Abduction vs. Induction	41
	Potential Answers and Informational Value	45
	Epistemic Utility	50
	Inferential Expansion	51
	Stable Acceptance and Caution	53
	Conclusion	56
3	CONTRACTION, REPLACEMENT, REVOLUTION, AND TRUTH	58
	Corrigibility	58
	Contraction	59
	Degrees of Corrigibility	61
	Contradiction vs. Anomaly	62
	Replacement	63
	Contextualism	65
	Revolutionary vs. Normal Science	67
	Myopia and Convergence to the Truth	70
	Theoretical and Practical	72

4	CREDENCE AND DELIBERATION	74
	The Four Questions	74
	Rational Credence	76
	Confirmational Commitments	79
	Bayes' Theorem and Conditionalization	83
	Inductive Logic	85
	Contextualism	91
	Rational Valuation	94
	Maximizing Expected Utility	95
	Revision	98
	Prospectus	100
5	CREDAL COHERENCE, CREDAL CONSISTENCY, AND <i>E</i>-ADMISSIBILITY	104
	Credence and Choice	104
	Classical Representations of Decision Problems	105
	Dominance	106
	Credal Coherence	108
	Weak Dominance and Credal Regularity	109
	Irrelevant Possibility and Conditional Probability	112
	Countable Additivity	118
	Credal Consistency	118
	On Infinitely Many Alternatives	119
	The Continuous Case	122
	Violating Countable Additivity	125
6	SUSPENDING JUDGMENT BETWEEN OPTIONS	132
	The Rule for Ties	132
	The Rule for Ties Extended	135
	<i>P</i> -Admissibility	137
	On the Strength of Options	139
7	SECURITY	144
	Decision Making Under Uncertainty	144
	Regret, Hope, and Security	147
	Spreads in the Odds	151
	On Appraising Security	156
	Mixed Options	162
8	VALUE CONFLICT	164
	Lack of Self-Awareness	164
	Unresolved Conflict	168
	Valuational Convexity	173
	Second Best and Second Worst	176
	Conflict vs. Ignorance	177
	Conflict in Epistemic Values	180

9	CREDAL CONVEXITY	183
	Varieties of Ignorance	183
	Black Boxes	186
	Metaphysical Probability	188
	Metaphysical Possibility	190
	Why Convexity?	191
	Convexity and the Multiplication Theorem	193
	Geometrical Representations	196
	Intervalism	197
	Measuring Credence	204
	Independence of Irrelevant Alternatives	208
	Supporting Lines and the Measurement of Credence	210
	Conclusion	214
10	CONDITIONALIZATION AND IRRELEVANCE	216
	Confirmational Commitments	216
	Conditional Probability	217
	Harper Functions	219
	Irrelevant Possibility vs. Serious Impossibility	222
	Confirmational Irrelevance	225
	Strong Confirmational Irrelevance	227
	The Finite Case	228
	Infinity	230
	Sufficiency and Irrelevance	230
11	ABILITIES AND DISPOSITIONS	234
	Chance and Credal Probability	234
	Ability and Serious Possibility	236
	Dispositions	237
	Disposition as Compulsion	238
	Explicit <i>N</i> -Predicates	239
	Ability Adverbially Modified	241
	Disposition as Primitive	243
	Knowledge of Ability and Appraisal of Serious Possibility	244
	Subjunctive Conditionals	246
	Sample Spaces	248
12	CHANCE AND DIRECT INFERENCE	250
	Chance and Credence	250
	Simple Chance Predicates	252
	Stochastic Irrelevance	253
	Direct Inference	254
	The Semantics of Chance	256
	Chance and Coherence	261
	Chance and Credal Consistency	261
	The "Existence" of Chance	264
	More on the Existence of Chance	267

	Chance and the Long Run	271
	Statistical Prediction	274
	Subjunctive Conditionals	276
	Composite Hypotheses: The Finite Case	278
	Composite Hypotheses: The Countably Infinite Case	280
	Composite Hypotheses: The Continuous Case	284
	On a Pseudo Paradox	284
13	INVERSE INFERENCE	288
	Direct Inference and Bayes' Theorem	288
	Objectivist Necessitarianism and the Relevance of Data	296
	Likelihood and Irrelevance	298
	Unbiased Priors	304
	Strongly Unbiased Priors	306
	Standardized Priors and Posteriors	308
	Stable Estimation	309
	Likelihood and Rejection Rules	311
	Ignorance and Strongly Unbiased Priors	315
	Prior Data	319
	Irrelevance Revisited	322
	Probability and the Growth of Knowledge	323
14	TAME FIDUCIAL INFERENCE	325
	Introduction	325
	Three Cases	326
	The Finite Case	328
	The Countably Infinite Case	335
	The Continuous Case	336
	Is Tame Fiducial Inference Convincing?	342
15	LIKELIHOOD	343
	Likelihood and Evidential Support	343
	The Law of Likelihood in the Discrete Case	347
	The Law of Likelihood and Fiducial Inference in the Finite Case	352
	The Law of Likelihood and Fiducial Inference in the Countably Infinite Case	353
	The Principle of Irrelevance in the Discrete Case	354
	The Law of Likelihood in the Continuous Case	355
	Likelihood and Fiducial Inference in the Continuous Case	358
	The Principle of Irrelevance in the Continuous Case	359
	The Law of Likelihood and Indifference	359
	Seidenfeld's Theorem	363
	Likelihood and Fiducial Probability: Another View	367

16	DIRECT INFERENCE AND CONDITIONALIZATION	369
	Fisher on Direct Inference	369
	Kyburg's Theory	375
	Kyburg's Theory and Fiducial Inference	382
	Kyburgian Direct Inference and Confirmational Conditionalization	383
	Dempsterian Conditionalization	385
	Random Selection vs. Random Membership	392
17	ON USING DATA AS INPUT	399
	Observation Reports	399
	Objectivist Necessitarianism a la Neyman-Pearson	403
	Security Levels	408
	Mixed Options	410
	Using One's Own Reports as Input	413
	Turning Evidence into Input	415
	The Long Run	419
18	THE CURSE OF FREGE	424
	APPENDIX	431
	A Brief Sermon on Assessing Accident Risks in U.S. Commercial Nuclear Power Plants	431
	REFERENCES	445
	INDEX	457

PREFACE AND ACKNOWLEDGMENTS

Knowledge is a resource for inquiry and deliberation. So understood it often stands in need of improvement. The enterprise of knowledge seeks to satisfy that need.

Even if cognitive resources render other services, at a very minimum they make an important contribution to the realization of the aims of persons and institutions regardless of whether the aims are political, ethical, economic, or scientific. We may, therefore, obtain some understanding of these resources through a grasp of the way they should function in rational decision making.

When justifying a choice between feasible options, a distinction needs to be made between hypotheses concerning consequences of these options which are possibly true and hypotheses concerning options which are not possible as far as the agent knows. To assess the expected values of these options a still finer discrimination among possibly true hypotheses with respect to probability is also helpful.

A body of knowledge serves in such deliberation as a standard for serious possibility. A credal state appraises the possibly true with respect to probability.

An account of the improvement of knowledge ought to explain in as systematic a fashion as is acceptable how revisions of knowledge and credence may be brought under critical control for the purpose of improving them as cognitive resources for deliberation and inquiry.

Contemporary preoccupation with the "incommensurability" of rival "paradigms" has tended to leave the impression that at critical points in the development of science we are compelled to make revisions in a manner which cannot be justified without begging questions. Such views are untenable both as they bear on the revision of corpora of knowledge and on the revision of credal states.

At the same time, I regard efforts to eviscerate inquiry by

removing human agents and institutions with their aims, problems, and resources from a centrally relevant role in an account of deliberation and inquiry to be an exaggerated response to the excesses of psychology, sociology, and historicism. We can attain objectivity enough without ascent to Popper's third world.

This book elaborates on these themes. The first three chapters discuss the functioning of a corpus of knowledge in inquiry and deliberation as a standard for serious possibility. The implication that, for the investigator, his corpus or standard should be regarded as infallibly true is explained and defended against the charge that it precludes reasoned revisions of knowledge. The correct view that knowledge is corrigible—i.e., subject to critically controlled revision—is severed from the thesis of fallibilism according to which knowledge is or ought to be possibly false from the agent's viewpoint.

Chapters 4–10 seek to extend the epistemological outlook already applied to the improvement of knowledge to probability judgment. The Bayesian ideal is rejected. According to that ideal, rational agents should always be committed to credal states of probability judgment representable by unique probability measures on the grounds that such an ideal prohibits suspension of judgment between rival numerically precise credal states relative to the same body of evidence. An ideally rational agent should sometimes suspend judgment by embracing a credal state which is indeterminate in the sense that more than one probability measure is considered permissible for the purpose of computing expected values in deliberation.

To be persuasive such an argument needs supplementation by an account of rational choice and statistical inference sufficiently comprehensive to cover cases where probability judgment is determinate and where it is indeterminate. The discussion of indeterminate probability judgment, indeterminate utility judgment, decision theory, inductive logic, chance, and statistical inference which takes up most of the space in this book seeks to establish the existence of such an account. The result is sufficiently comprehensive, so I hope, to recognize Bayesian prescriptions as legitimate in special cases while acknowledging that Bayesian doctrine is not as universally applicable as it claims to be. Yet, the aim is to avoid the

skepticism and eclecticism so characteristic of anti-Bayesian critics by supplying a rival to Bayesianism pretending to the scope so many Bayesians claim for their doctrine.

I learned from reading R. B. Braithwaite's *Scientific Explanation* that no useful semantics for predicates attributing objective probabilities or chances to systems can be given; and that clarity concerning the connections between chances and test behavior can be obtained only through studying direct inference from knowledge of chances to judgments about test behavior and inverse inference from knowledge of test behavior to judgments about chances.

Building on the results of the first ten chapters, chapters 11–13 offer an account of direct and inverse inference linking chance with test behavior. Chapters 14–18 contrast my account with other alternatives.

My original intent was to extend and improve upon the epistemological outlook presented in my book, *Gambling with Truth*. In addition to contributing to the elaboration and modification of my own epistemological outlook, however, the analytical apparatus I have constructed should interest students of epistemology, philosophy of science, decision theory, probability theory, and statistical inference even if they are not committed and are, indeed, opposed to my orientation.

Ever since reading about "spans of confirmation" in Frederic Schick's doctoral dissertation, I have been intrigued with the idea of representing probability judgments by sets of probability distributions. Over the years we have argued, often vehemently, about topics in value theory, individual decision making, collective decision making, and probability theory. These discussions have been invaluable to me and have touched every aspect of the account of rational choice offered here.

In earlier days, I reacted to Henry Kyburg's worries about acceptance. His pioneering work on indeterminate probability has been one of the sources of my ideas for this volume. But I am especially grateful to him for the opportunity to disagree with his views on direct inference. In disagreeing with him, I have learned more about this topic than from anyone else.

I initially thought the originality of the scheme I was developing for using indeterminate probabilities (as compared with the pioneering work of B. O. Koopman, Kyburg, I. J.

Good, C. A. B. Smith, and Schick) resided chiefly in the account of inductive probability, logical probability, and decision theory and in the account of the revision of probability judgment. Howard Stein helped me to see that the technical articulation of my ideas on probability judgment departed significantly from ordinalist and intervalist views which currently dominate discussions of indeterminate probability.

Teddy Seidenfeld was a graduate student at Columbia while this book was being written. It has been my good fortune to have supervised and to have learned from his doctoral dissertation. His insightful work on the fiducial argument has been of great help to me in writing chapters 14-17 and traces of my discussions with him appear throughout this book.

This book was written at a time when interest in modal logic and its applications to metaphysics and epistemology dominated much discussion. I have been fortunate to have had the opportunity to discuss some of my reservations with James Higginbotham and Charles Parsons. I remain convinced that this new wave is a retrograde step in philosophy but, thanks to Higginbotham and Parsons, my reservations are more informed than they otherwise would have been.

I regard this present work as a fragment of an account of how the context of inquiry directs the revision of knowledge. This was Dewey's problem. I became interested in it under the tutelage of Ernest Nagel and Sidney Morgenbesser.

I was prompted to check into the methodological discussions in the Rasmussen report thanks to a conversation with Seymour Melman and Raymond Siever. This led to the belated addition of the appendix to the book.

Henry Kyburg, Sidney Morgenbesser, Ernest Nagel, Charles Parsons, Frederic Schick, Teddy Seidenfeld, Stephen Spielman, and Howard Stein have read various portions of this book at some stage in its composition and have given me their critical reactions. I wish to thank them all.

Work on this book began during the winter of 1973 while I was a visiting scholar at Corpus Christi College, Cambridge. I wish to thank the Master and Fellows of Corpus Christi and the members of the Faculties of Philosophy and History and Philosophy of Science at Cambridge for their gracious hospitality which facilitated my embarking on this project. Partial

support for this work was supplied by the National Science Foundation.

The writing of this book has been long and arduous. At various stages, my wife, Judith, has helped me remove important infelicities from the prose and avoid needless digressions. More importantly, she has offered me the love and support I required to see the project through. I dedicate this volume to her.

Isaac Levi
New York City
October, 1978

**1.1
Knowledge
without
Pedigree**

Knowledge is widely taken to be a matter of pedigree. To qualify as knowledge, beliefs must be both true and justified. Sometimes justification is alleged to require tracing of the biological, psychological, or social causes of belief to legitimating sources. Another view denies that causal antecedents are crucial. Beliefs become knowledge only if they can be derived from impeccable first premises according to equally noble first principles. But whether pedigree is traced to origins or fundamental reasons, centuries of criticism suggest that our beliefs are born on the wrong side of the blanket. There are no immaculate preconceptions.

Where all origins are dark, preoccupation with pedigree is self-defeating. We ought to look forward rather than backward and avoid fixation on origins.

Epistemologists should heed similar advice. Whatever its origins, human knowledge is subject to change. In scientific inquiry, men seek to change it for the better. Epistemologists ought to care for the improvement of knowledge rather than its pedigree. They ought to ask what *X* (who may be a person or a group) should do, given his knowledge at a time *t*, to render that knowledge more efficient in performing its functions.

Even though attention is shifted from the pedigree of knowledge to its improvement, a question of justification, nonetheless, remains. *X* should not modify his body of knowledge unless in doing so he improves it. Hence, even though *X* need not justify having *h* in his corpus of knowledge once he has accepted it, prior to doing so he will be under some obligation to justify adding *h* to his body of knowledge. Furthermore, should he contemplate removing *h* from his corpus once it is in, he ought to consider whether such a deletion is warranted.

Prior to adding *h*, *h* is, from *X*'s point of view, a hypothesis being entertained for naturalization into *X*'s corpus of knowledge. Once *X* has concluded that adding *h* to his corpus is justified as an improvement to his corpus and has implemented

his decision, *h* ceases, for *X*, to be a hypothesis. It has become a premise, evidence, settled assumption, or part of the "background knowledge" to be used in subsequent inquiries into the credentials of other statements as well as in practical deliberations aimed at moral, political, economic, or other practical objectives. Whether *h* is a theory, law, statistical claim, or observation report, and regardless of the grounds on which it has been added, its status as an item in *X*'s corpus has been settled and the grounds on which it has been added no longer matter.

To be sure, the status of *X*'s belief that *h* as knowledge is by no means secure—even from *X*'s point of view. *X* can consistently recognize that occasions might arise where his body of knowledge will be improved by removing *h*. But just as adding new items to a corpus of knowledge requires justification, so does removing an item. In addition, the justification should be based on the assumptions in *X*'s corpus prior to removing *h*—i.e., when *h* is still in the corpus and is still entitled to the status of knowledge. In general, neither the origins of *X*'s having *h* in his corpus nor the grounds on which he justified adding *h* to his corpus in the first place will be relevant to deciding whether to remove *h* unless consideration of origins can be shown to have a bearing on whether elimination of *h* will improve *X*'s corpus.†

**1.2
Knowledge as
a Standard for
Serious Possi-
bility**

X's knowledge at time *t* is a resource he uses in subsequent inquiries and deliberations. What are the uses to which *X* puts his knowledge or is committed to putting his knowledge?

† I have been talking here of *X*'s justifying a revision of his corpus to himself. There are, of course, other contexts calling for justification. *X* may be concerned to justify *X*'s revising his corpus in a certain manner to *Y*. A more important task would be for *X* to justify to *Y* *Y*'s revising *Y*'s corpus in a certain manner. In this last case, *X* might, for example, already accept *h* in his own corpus. In his effort to show *Y* that *Y* should add *h* to his corpus, *X* should not appeal to his (*X*'s) corpus containing *h* but to a corpus of "shared agreements" with *Y*. In doing this, *X* is in no way acknowledging that his (*X*'s) acceptance of *h* stands in need of justification to *X*. *X* is not in doubt as to the truth of *h*. *Y* is. That is the point of *X*'s effort to show *Y* that he should add *h* to his corpus.

In spite of the pragmatic differences in these several sorts of justification, I contend that justification always involves justification of a shift from one corpus to another where the contents of the initial corpus are the "premises" of the justification.

X's knowledge at *t* serves as a standard for distinguishing truth-value-bearing hypotheses whose truth is a serious possibility according to *X* at *t* from those whose truth is not a serious possibility according to *X* at *t*.

When witnessing the toss of a coin, *X* will normally envisage as possibly true the hypotheses that the coin will land heads up and that it will land tails up. He may also envisage other possibilities—e.g., its landing on its edge. However, if he takes for granted even the crudest folklore of modern physics, he will rule out as impossible the coin's moving upward to outer space in the direction of Alpha Centauri. He will also rule out the hypothesis that the Earth will explode.

Of course, it is logically possible that the coin will fly out towards Alpha Centauri upon tossing and even that the Earth will explode under these conditions. But this means only that if *X* were in that (different) cognitive state wherein his standard for serious possibility consists exclusively of logical truths, these hypotheses would be serious possibilities according to *X*.†

In real life, men do not restrict their knowledge to logical truths. If *X* were offered a gamble on the outcome of the toss of the coin, he would not take seriously the logical possibility that the coin will fly out towards Alpha Centauri. Nor would he take into account the logical possibility that the Earth will explode. Logically possible though these hypotheses may be, *X* both would and should ignore them because they are not seriously possible from his point of view. His knowledge rules them out.

Judgments of subjective or credal probability are intimately related to evaluations of hypotheses with respect to serious possibility. If hypothesis *h* bears positive credal probability according to *X* at *t*, *X* should evaluate the truth of *h* as a serious possibility. If it is not a serious possibility according

† In the following discussion, I shall regard rational agents as prohibited from adopting the set of logical truths as their standard of serious possibility. The weakest such standard allowable shall be the urcorpus *UK* discussed in section 1.3. I shall call the truth of *h* logically possible if and only if *h* is consistent with *UK*. If *h* is logically possible in this sense, it is, of course, logically possible in the usual sense; but the converse fails. This practice suppresses the distinction between logical and mathematical possibility which may loom large in other contexts but which will rarely be of any importance in this discussion.

to X at t that h and g bear different truth values, X should assign h and g the same credal probability at t .

We should not conclude that if h bears 0 credal probability according to X at t that h is not a serious possibility according to X at t . X might assume that a rod has a length representable by a real number in some interval without knowing what that real number is. According to X , noncountably many rival hypotheses are serious possibilities. Yet, he might assign them all 0 credal probability.

Care should be taken to distinguish hypotheses discounted in practical deliberations because they are not serious possibilities from hypotheses which are ignored because their truth values are irrelevant given the goals and values being promoted in deliberation. When facing the decision as to whether to accept the bet, it may be a serious possibility according to X that the coin will land heads and a Republican president will be elected in 1980. It will also be a serious possibility that the coin will land heads and a Republican will not be elected. But X will ignore these alternatives. Instead, he will focus on whether the coin will land heads regardless of who will be president in 1980. That information has no importance for him insofar as he is concerned to decide whether to accept or reject the bet. All relevant possibilities must be serious ones; but the converse need not hold.†

† Failure to attend to the distinction between serious and relevant possibility contributes to legitimizing some pernicious forms of wishful thinking.

Once X 's corpus at t is fixed, the distinction between what is and what is not a serious possibility according to X at t should be fixed. Variation in the problems X addresses and the goals and values he seeks to promote should not justify alteration of the distinction as long as the corpus remains fixed.

But even if the corpus remains fixed, a change in X 's problems, goals, and values could alter his assessment of which serious possibilities are relevant possibilities to be taken into account in X 's deliberations.

Consequently, if care is not taken and relevant possibility is confused with serious possibility so that assigning h the status of a hypothesis which is not relevantly possible suffices to mandate assigning it 0 credal, personal, or subjective probability, a change in X 's values could justify a revision in X 's state of probability judgment. This is the pernicious form of wishful thinking to which I am referring.

Thus, Popper writes: "What is the upshot of all this? It is that 'absolute certainty' is a limiting idea, and that experienced or subjective 'certainty' depends not merely upon degrees of belief and upon evidence, but also upon the situation—upon the importance of what is at stake." (*Objective Knowledge*. London: Oxford University Press, 1972, p. 79) Perhaps "wishful thinking" is inaccurate. Someone is guilty of wishful thinking if he tends to increase

The distinction between logical possibilities which are serious and those which are not is important in scientific inquiry as well as practical deliberation. When devising hypotheses about the constitution of quasars, no one considers the logical possibility that they are conglomerations of drosophila flies to be a serious possibility. In designing experiments, one takes into account only those hypotheses as to the outcome which are serious possibilities.

Nothing I have said or will say amounts to an explication of the concept of serious possibility. That is to say, no definition in other terms will be offered. I have indicated that "serious possibility" is relational: h is a serious possibility according to X at t . Moreover, a necessary condition for X to assign positive credal probability to h at t and to do so legitimately is that the truth of h be a serious possibility according to X at t . In later chapters, an account of rational choice and scientific inquiry will be outlined which will furnish more structure to the concept of serious possibility by articulating its relevance to deliberation and inquiry.

Assuming, however, some understanding of the notion of serious possibility, I am advancing a thesis about the use of knowledge in deliberation and inquiry—to wit, that X 's body of knowledge, evidence, or settled assumptions at t is his standard for serious possibility at t . h is a serious possibility according to X at t if and only if h is consistent with his corpus of knowledge at t .

1.3 Incorrigibility

If X 's corpus of knowledge at t serves as X 's standard for serious possibility at that time and if, in addition, X 's corpus of knowledge is subject to revision, X 's standard for serious possibility is also subject to change.

But if we are to concentrate on such revision in a systematic manner, we need some way of representing a corpus of knowledge or standard for serious possibility so that we can mark changes in corpus or standard.

his judgments of probability in favor of hypotheses he wants to be true. Popper seems to think that the more attractive a hypothesis and the more important the issue at stake, the less probable it should be judged to be. I cannot see, however, that Popper's prescription is any less pernicious than a recommendation to think wishfully. Both approaches see judgments of probability as depending on the agent's values and desires.

To this end, I propose that we focus attention on X 's standard for serious possibility for hypotheses expressible in a suitably regimented language L . I do not suppose that X speaks or reads L or is in any other way familiar with it. I use L as a device for discussing revisions in standards for serious possibility. At least some revisions, so I shall suppose, are representable as changes in the way X is committed to evaluating hypotheses expressible in L with respect to serious possibility; and those which are not so representable can be reached (again so I assume) by enriching L in appropriate ways.

X 's standard for serious possibility, insofar as it is expressible in L , is a set $K_{N,t}$ of sentences in L such that all hypotheses expressible in L which are seriously possible according to X at t are consistent with $K_{N,t}$, and only such hypotheses are.

Throughout this book I shall be concerned only with those revisions of knowledge which are tantamount to revisions in X 's standard for serious possibility. This decision on my part has some implications for how a corpus expressible in L is to be understood.

Suppose that $\sim h$ is inconsistent with X 's corpus $K_{N,t}$ expressible in L . Yet, h , though deducible from $K_{N,t}$, is not a member of $K_{N,t}$. We could then consider the addition of h to $K_{N,t}$ as a change in knowledge; and there are important contexts where it is useful to mark such changes. Thus, X might prove a new theorem or, after having taken a course in some branch of mathematics, discover that h is a consequence of assumptions he already has made.

Such changes, however, are not changes in X 's standard for serious possibility or, at any rate, in the evaluations of hypotheses in L with respect to serious possibility which X is committed to making by his standard. But it is only changes in the evaluations of hypotheses in L with respect to serious possibility which X is committed to making by his standard that are of concern to me in this investigation. Consequently, changes of the sort just described are not at the center of attention.

By the same token, I do not propose to consider any sort of revision in the standards for serious possibility which allows

hypotheses inconsistent with truths of first-order logic, set theory, or mathematics to be serious possibilities.

Both of these decisions impose limitations on the scope of my account of the revision of knowledge. I am concerned with changes in knowledge which are changes in standards for serious possibility and I do not wish to countenance changes which allow the falsity of truths of first-order logic, set theory, or mathematics to be serious possibilities.

These considerations suggest that we can, for the purposes of this essay, represent a corpus of knowledge expressible in L by a deductively closed set of sentences in L containing a set UK (to be called the *urcorpus*). UK consists of all logical truths, set theoretical truths, and mathematical truths expressible in L , as well as any other items expressible in L which might qualify as incorrigible in the sense that they are immune from removal from the standard for serious possibility (or are not possibly false according to any allowable standard).

On what basis should a hypothesis be considered eligible for membership in the urcorpus of incorrigible claims UK ? My aim is to see into UK those assumptions which any corpus should have if an account of the revision of knowledge of the sort I seek to construct is to stand a chance of working. But I do not pretend that I have shown that the items I include in the urcorpus meet this (vaguely formulated) standard.

Categorical incorrigibility should be distinguished from idiosyncratic incorrigibility. X (whether X is a person, or a group, or an institution) may need to presuppose X 's existence by some version of the *cogito* argument, although he does not need to assume that Y exists as an incorrigible assumption.

The distinction between categorically incorrigible assumptions in an urcorpus and sentences which, if they are in X 's corpus at t , are liable to removal under suitable circumstances does not coincide with distinctions between analytic and synthetic truths, a priori and a posteriori truths, or conceptual and nonconceptual truths. Conceptual, a priori, or analytic truths are often taken to be open to revision. However, it is typically maintained that the sort of revision of knowledge allowed is a revision in conceptual framework and is to be distinguished from the sort of changes with which I am interested here. In particular, as I shall explain shortly, avoidance of errors should be a desideratum of all legitimate revisions of

knowledge. According to those who allow for changes in conceptual framework, such changes entail a change in the very conception of error and, hence, cannot be evaluated with respect to prospects of importing or avoiding error from the vantage point of the conceptual framework employed prior to change in any intelligible and non-question-begging manner.

In effect, I have already conceded that changes in logic or set theory are to count as changes in conceptual framework in this latter sense. But I do not propose to discuss such changes. As for the rest, I propose to outline an account of the revision of knowledge which minimizes the significance of conceptual changes. In particular, I wish to emphasize the importance of avoidance of error as a desideratum in inquiry and to discount the excuse for ignoring this desideratum which appeals to the claim that a certain sort of change is a conceptual change.

1.4 Truth

Setting to one side those items I have identified as categorically or idiosyncratically incorrigible, I wish to maintain that all items in X 's corpus $K_{X,t}$ are not merely open to revision but are corrigible in the sense that avoidance of error should be taken into account in revising the corpus. That is to say, X should take avoidance of error into account in contemplating alternative revisions of his corpus.

But since X 's corpus at t is his standard for serious possibility at t , then, by his lights, no item in that corpus is possibly false. Hence, they are all true. My contention is that X should be concerned to avoid error in a sense that presupposes the truth of all items in his corpus.

Insofar as we represent X 's corpus at t by sentences in L , the conception of truth involved may be characterized by a definition of "true in L ." Once more, I do not assume that X is consciously or explicitly aware of such a truth definition or even that he has any mastery of the language L . The truth definition is intended to represent a feature of his goals in seeking to revise his knowledge.

I suppose that the definition of "true in L " proceeds along Tarskian lines.

A hierarchy of more and more inclusive languages $L, L_1, \dots, L_i, \dots$ is introduced. For each L_i , there is an appropriate urcorpus UK^i . Each UK^i contains assumptions sufficient to

provide a minimal definition of "true in L_{i-1} " along Tarskian lines. However, as just noted, if X 's corpus expressible in L_i at t is $K_{X,t}^i$, the definition of "true in L_{i-1} " which captures the sense of truth and, hence, of error in which X is concerned to avoid error when revising $K_{X,t}^i$ is one which assumes as part of the truth definition that all items in $K_{X,t}^{i-1}$ are true in L_{i-1} .

1.5 Comparative Statics

In section 1.3, I introduced the requirement that every X adopt as his corpus expressible in L a body of assumptions representable by a deductively closed set of sentences in L containing the urcorpus UK for L . In section 1.4, I implicitly suggested extending that representation so that it becomes a sequence of sets of sentences $K_{X,t}, K_{X,t}^1, \dots, K_{X,t}^i, \dots$ expressible in languages $L, L_1, \dots, L_i, \dots$, respectively, meeting the following conditions:

- (a) each $K_{X,t}^i$ contains the urcorpus UK^i for the appropriate L_i ,
- (b) $L_i \subseteq L_{i+1}$ and $K_{X,t}^i \subseteq K_{X,t}^{i+1}$,
- (c) $K_{X,t}^{i+1}$ furnishes a characterization of "true in L_i " meeting Tarskian requirements and implying the truth of all sentences in $K_{X,t}^i$.

This complication in the mode of representation should not loom large in subsequent discussions; for insofar as we are focusing attention primarily on changes in knowledge expressible in L , we need attend to changes in corpora at higher levels only insofar as they are automatically induced by the changes in the corpus in L .

On the other hand, the introduction of the hierarchy returns us to an issue only partially considered in section 1.3.

Suppose that X is interested in the integer in the billionth decimal place in the decimal expansion of π . Of the ten hypotheses of the form "The integer in the billionth decimal place in the decimal expansion of π is j " where j is a standard designator for one of the first 10 nonnegative integers, exactly one is consistent with the urcorpus UK which, it should be remembered, contains mathematical truths. On the other hand, since we may suppose that X has not made the required calculations and cannot resort to a calculator of any sort for help, there is an important sense in which X does not know which of these hypotheses is entailed by UK and, hence, does

not know what the integer in the billionth decimal place of the decimal expansion of π is.

In section 1.3, I contended that the sort of change which would take place were X somehow to make the required calculations is not a change in X 's standard for serious possibility and, hence, could be ignored in this investigation.

But it may be objected that when X fails to make the calculations he is in no position to rule all but one of the ten hypotheses out of consideration as serious possibilities. Consequently, making the calculations does lead to a change in X 's standard for serious possibility.

At this point, I suggest that we distinguish between the standard for serious possibility to which X is committed at time t and X 's awareness at t of the standard to which he is committed; or equivalently of the corpus to which he is committed. In our example, X is not aware of all his commitments due to an inability to make certain calculations. On other occasions, X 's memory may fail him. Or X may suffer from emotional disturbance (if X is a person) or some social disturbance (if X is an institution).

If X is committed to a standard for serious possibility, he should live up to that commitment to the extent that he is able. Such ability depends, in part, on the extent to which X is aware of that commitment. Men and societies vary widely in their capacities for making computations, for storing information, and for maintaining emotional and social stability. A normative account of the improvement of knowledge should prescribe no more than persons and institutions are capable of implementing. Thus, it would be foolish to require that rational X identify all the logical consequences of the assumptions he explicitly makes. We can, at most, expect him to identify those consequences insofar as he is able.

The standard for serious possibility to which X is committed at t is, on this view, the standard to which he would be conforming were he ideally situated (i.e., endowed with perfect computational facility, memory, and emotional or social health) and were he also rational. I do not suppose, however, that agents (persons or institutions) are ideally situated and rational. I do take them to be real agents with commitments of various kinds—including commitments to standards for se-

rious possibility; and I urge them to be rational in the sense that they live up to their commitments insofar as they are able.

Thus, changes in X 's awareness of his commitments at t ought to be distinguished from changes in X 's commitments. The former sort of change may be compared to a shift towards an equilibrium. The more fully aware X is of his commitments, the closer he is to a state of cognitive equilibrium. On this analogy, X is committed at t to a state of cognitive equilibrium whether he has actually attained it or not. The features of rational equilibrium I have been discussing have been introduced by an appeal to those functions which X 's corpus of knowledge ought ideally to perform.

Given an adequate characterization of states of cognitive equilibrium, the account of the revision of knowledge built on it may be viewed as an analogy to what is sometimes called "comparative statics." Thermodynamics and some branches of economic theory illustrate comparative static theories which investigate changes in equilibrium states of systems suitably specified without scrutinizing the details of the paths such systems follow in moving from one equilibrium state to another. The normative analogue of such theories of the sort I am aiming to construct here prescribes shifts from one state of cognitive equilibrium to another without prescribing details of the psychological or social changes which are made in implementing the revision.

Thus, although on my approach a change in cognitive state will be represented by a shift from one deductively closed set of sentences in L to another, I am not recommending that rational agents immediately make revisions in this way. I do not assume that men or institutions are able to realize equilibrium states. I do believe they are able to approach equilibrium to a sufficiently good approximation to render it less than utterly quixotic to regard them as committed to making the effort. At the same time, on the comparative static approach, one is absolved from prescribing the details of the psychological or social processes to be undergone in attempting to approach equilibrium.

For these reasons, the only relevant way in which psychology or sociology can be used to criticize the prescriptions concerning changes in cognitive equilibrium of interest here is by showing that men or institutions are incapable of moving

toward equilibrium or making changes in commitments by some path or other. It is not enough to show that men often fail to conform to the dictates of such norms. And since the norms do not prescribe the details of the social or psychological paths to be followed in changing commitments, they cannot be rejected on the grounds that they cannot be implemented by following some specific path or because they cannot be implemented in a certain fixed period of time. Psychology, the social sciences, and various technologies (such as computer technology) are of relevance to the topics of interest here insofar as they can contribute to our ability to live up to our commitments by improving our computational facility, memory, and emotional or social health. When they focus instead on the extent of our incapacities and disabilities without seeking remedies, they are not doing the job that is needed.

If X fails to live up to his commitments or is unaware of what these commitments are, his state of knowledge is in cognitive disequilibrium. I lack an adequate analysis of awareness and have no criterion of awareness in terms of linguistic or other behavior or in terms of introspection of "internal" states. I do assume, however, that to be aware of one's commitments is to know what they are.

Hence, in a state of cognitive equilibrium, X should know which hypotheses (expressible in L) are possibly false and which are not possibly false according to his mode of evaluating hypotheses with respect to serious possibility at that time. This condition of self-knowledge is expressed by the following requirements to be added to conditions (a)-(c) cited previously:

- (d) if $h \in L_t$ and $h \in K_{X,t}^1$, $\{h \in K_{X,t}^1\} \in K_{X,t}^{1+1}$;
- (e) if $h \in L_t$ and $h \notin K_{X,t}^1$, $\{h \notin K_{X,t}^1\} \in K_{X,t}^{1+1}$.

For most purposes, we may ignore the hierarchy of corpora expressible in the sequence of metalanguages which represents X 's state of knowledge in cognitive equilibrium at t . We need consider the corpus expressible in L alone. We shall wish to consider, to be sure, the alternative corpora expressible in L to which X could become committed under some circumstances or other. Such *potential corpora* expressible in L are deductively closed sets containing UK .

UK represents the corpus of all incorrigible assumptions, as

explained previously. It may be construed as representing the cognitive equilibrium state of an agent in a state of extreme modal ignorance.

Potential corpora in L are partially ordered by the set-inclusion relation. K is weaker than K' if and only if $K \subset K'$. The weakest corpus in the lattice so generated is, of course, the uncorpus UK . The strongest is the inconsistent corpus.

1.6 Infallibilism

The central theme of the preceding sections has been that X 's corpus of knowledge or evidence at t serves as his standard at t for discriminating between those hypotheses whose truth is a serious possibility according to X at t and those whose truth is not a serious possibility. This theme has motivated the way in which a body of knowledge is formally represented and how a change in knowledge is to be understood. Above all, it supplies a key element in understanding how X 's corpus of knowledge functions or should function in practical deliberation and scientific inquiry.

An immediate consequence of the thesis that X 's corpus at t serves as his standard for serious possibility at t is that, according to X at t , no item in his corpus at t is possibly false in the sense of serious possibility. If $h \in K_{X,t}$, then h is infallibly true according to X at t in a straightforward and important sense. Thus if Y should disagree with X at t , from X 's point of view, Y is certainly in error. It would be inconsistent for X to concede to Y that he (X) might be mistaken if by this X acknowledges the falsity of h as a serious possibility.

Of course, X may consistently acknowledge that items he accepted in his corpus at previous times or will assume in the future are possibly false. He is not committed to the view that whatever he has endorsed in the past or will accept as evidence in the future is infallibly true. But at t , X is committed to the view that whatever he assumes as part of his corpus at t is infallibly true.

Denying this thesis of *epistemological infallibilism* entails rejection of the view of knowledge as a standard for serious possibility and, hence, the view of how knowledge functions as a resource for inquiry and deliberation which I am advocating. Those who reject epistemological infallibilism, therefore, are under some obligation to supply an alternative view of the functions and value of knowledge.

Perhaps, however, epistemological infallibilism is not at odds with the fallibilist epistemologies so widely advocated by modern authors.

Charles Peirce characterized fallibilism as the doctrine "that we can never be sure of anything"¹ or "that we cannot attain absolute certainty concerning matters of fact."² Peirce undoubtedly intended to deny that men can attain *permanent* certainty of matters of fact. But this may mean either that all logical possibilities (or, more accurately, items consistent with *UK*) are or ought to be serious possibilities (so that rational agents should never attain certainty of matters of fact) or that standards for serious possibility are subject to revision (so that *X* may be certain at *t* that *h* is true and be justified in ceasing being convinced at *t'*).

This claim is quite compatible with epistemological infallibilism and the claim that knowledge is used as a standard for serious possibility. Rejection of permanent certainty, together with the thesis that knowledge is the standard for serious possibility, implies either that *X* should stick permanently to the urcorpus as his corpus of knowledge or that knowledge is corrigible.

Peirce, however, seems to have intended to claim more than that we cannot attain permanent certainty concerning any matter of fact. He meant to deny that we can attain *maximum* certainty—at least when we adopt the scientific attitude.† According to Peirce's fallibilism, at no time should *X* discount the falsity of any matter of fact as not a serious possibility. Thus all logical possibilities (i.e., hypotheses consistent with the urcorpus *UK*) should be considered serious possibilities by every *X* and at every time. I shall call this thesis *categorical fallibilism*.

Categorical fallibilism is consistent with epistemological infallibilism. When the two theses are conjoined, rational *X* is

† C. S. Peirce, *Collected Papers* (Cambridge, Mass.: Harvard University Press, 1931), v. 1, p. 347. This passage supports a conclusion concerning Peirce's intentions which seems obvious in any case—to wit, that scientific men should not be maximally certain about any matters of fact. Since, in my opinion, Peirce failed to distinguish between maximum and permanent certainty when discussing fallibilism, many of his writings reveal, to anyone who is sensitive to the distinction, considerable ambiguity. Hence, although I believe my interpretation of Peirce is a fair one, it is difficult to offer citations so free from ambiguity as to preclude alternative readings.

obliged to restrict his corpus of knowledge to the urcorpus *UK*. Thus, the conjunction of the thesis that knowledge serves as the standard for serious possibility in inquiry and deliberation and the doctrine of categorical fallibilism implies the *in corrigibility* of knowledge.

Here we come to the crux of the matter. For Peirce, a central feature of the doctrine of fallibilism is its commitment to the corrigibility of human knowledge. "The scientific spirit requires a man to be at all times ready to dump his whole cartload of beliefs, the moment experience is against them."³ The thesis of corrigibilism, however, contradicts the joint assertion of categorical fallibilism and the claim that knowledge is the standard for serious possibility.

One alternative is to give up the thesis of corrigibilism and restrict knowledge to the contents of the urcorpus. Some empiricists modify this rather bleak view by allowing the urcorpus *UK* to be augmented by observation reports. I agree with Peirce in refusing to follow either of these courses. Predictions about the future as well as conclusions about the past may find their way into *X*'s corpus. So may theories, laws, and statistical claims. Of course, corpora of this sort are subject to critical review and revision. The thesis of corrigibilism ought, therefore, to be endorsed.

We are faced, therefore, with a choice of abandoning the thesis that knowledge is a standard for serious possibility—and thus also its corollary, the thesis of epistemological infallibilism—or of rejecting the thesis of categorical fallibilism.

I favor abandoning categorical fallibilism while endorsing epistemological infallibilism and knowledge as a standard for serious possibility.

Peirce took the opposite view:

We have seen how success in mathematics would necessarily create a confidence altogether unfounded in man's power of eliciting truth by inward meditation without any aid from experience. Both its confidence in what is within and the absolute certainty of its conclusions lead to the confusion of *a priori* reason with conscience. For conscience, also, refuses to submit its dicta to experiment, and makes an absolute dual distinction between right and wrong. One result of this is that men begin to rationalize about questions of purity and integrity, which in the long run, through moral decay, is unfavourable to science. But what is worse, from our point of view,

they begin to look upon science as a guide to conduct, that is, no longer as pure science but as an instrument for a practical end. One result of this is that all probable reasoning is despised. If a proposition is to be applied to action, it has to be embraced or believed without reservation. There is no room for doubt, which can only paralyze action. But the scientific spirit requires a man to be at all times ready to dump his whole cartload of beliefs, the moment experience is against them. The desire to learn forbids him to be perfectly cocksure that he knows already. . . . Thus the real character of science is destroyed as soon as it is made an adjunct to conduct; and especially all progress in the inductive sciences is brought to a standstill.⁴

Peirce's contrast between the importance of probabilistic reasoning in science and its unimportance in practical conduct is surely mistaken. In practical decision making, account should often be taken of risks and, hence, of probabilistic considerations.

Peirce is nearer to the mark when he concedes that in practical deliberation some logical possibilities are discounted as not being serious possibilities. Peirce insists, however, that in scientific inquiry all logical possibilities are serious.

According to Peirce, therefore, the standard for serious possibility used in science cannot be the same as that used in practical deliberation. To the extent that science supplies us with extralogical knowledge, that knowledge cannot serve as a standard for serious possibility in science (where all logical possibilities are serious) or in practical deliberation; for, so Peirce maintains, science should not be looked upon as a guide to conduct.

Thus, Peirce seems to reject two assumptions I am making: the thesis that knowledge serves as a standard for serious possibility (which entails epistemological infallibilism) and the thesis that rational *X* should, during any minimal interval of time, be committed to a single standard for serious possibility both for theoretical inquiry and for practical deliberation. In these two respects, Peirce drives a wedge between theory and practice I seek to remove.

To be sure, there is a measure of truth in the double-standard thesis. Recall that an agent *X* need not be a person. *X* can be a community such as a scientific community or a political party or a school or some institution which has goals,

makes decisions, and adopts various propositional attitudes. I follow the practice of students of the theory of consumer demand who attribute demand curves and even utilities to consumers even though the consumer may be a family or a firm and not a person. Obviously, to imitate this practice entails sweeping many problems under the rug. I am concerned, however, with an account of revisions of standards for serious possibility applicable regardless of who or what the agent is, provided the agent can adopt propositional attitudes, make decisions, and engage in inquiry and deliberation.

Observe, however, that one and the same *person* may belong to different communities each of which has its own goals and values and each of which may have different and sometimes conflicting standards for serious possibility.

It is at least entertainable that when *X* identifies himself as a member of community *Y*, he commits himself to the standard for serious possibility adopted by that community. But *X* might also identify himself with community *Z* and in the same way. This need not be the case, but could be so. Insofar as *X* finds himself in this predicament, he may be committed to a cognitive schizophrenia embracing different standards for serious possibility corresponding to his diverse roles. Though *X* is a single person, he may be several agents.

There is another way to describe the situation. We may say that *X* is a single agent with a single inconsistent standard for serious possibility derived from his joint commitment to conflicting standards for serious possibility endorsed by diverse communities.

I do not care whether *X* is described as a single person who is several agents each committed to a single standard for serious possibility or a single agent with an inconsistent standard. The kernel of truth in Peirce's double standard view is that owing to the fact that persons belong to different communities they will typically exhibit schizophrenia or inconsistency.

Peirce seems to accept this circumstance with equanimity. It is right and proper that the man of science who is also a member of other communities will be committed to different and conflicting standards for serious possibility at a given time. He should not take steps to remove the conflict even if he is capable of doing so. In science, every logical possibility should

be a serious possibility. In practical affairs it should not. There is no way in which *X* can gratify the demands of science and practice without retaining the conflict. The gulf between theory and practice entails either schizophrenia or inconsistency.

In rejecting the double standard, I stand opposed to such equanimity. It is no doubt neither psychologically nor socially feasible to remove all inconsistency or schizophrenia. But insofar as it is feasible to do so, all agents, whether they are individual persons or communities, should commit themselves to single consistent standards for serious possibility. When inconsistency arises, as it inevitably will, and when it is detected, steps should be taken to eliminate it.

But once the double standard is rejected, we must either insist that both in science and in practice all logical possibilities are serious or ought to be, or we should abandon categorical fallibilism both in science and in practical deliberation. I favor the latter alternative.

Once categorical fallibilism is abandoned, both in science and in practical deliberation, there is no contradiction in accepting both the thesis that knowledge should be the standard for serious possibility and the thesis of the corrigibility of knowledge.

The resulting view is infallibilist in two special senses: (a) The thesis that knowledge is the standard for serious possibility implies the thesis of epistemological infallibilism. From *X*'s point of view at *t*, there is no serious possibility that any item in his corpus at *t* is false. (b) Rejection of categorical fallibilism implies that standards for serious possibility should sometimes admit the impossibility of many extralogical (extra-set-theoretical and extramathematical) propositions.

On the other hand, my position remains fallibilist insofar as fallibilism is equivalent to corrigibilism—the thesis that *X*'s standard for serious possibility (or corpus of knowledge) is sometimes legitimately subject to revision. On this score, there is no disagreement between my view and Peirce's or, for that matter, Karl Popper's (who, with respect to the issues being considered here, appears to be in agreement with Peirce).

In advocating categorical and epistemological infallibilism, I also do not intend to claim infallibility for any source of

information such as the Delphic oracle or the Pope where a source of information is infallible if whatever it reports is true.

Finally, even *X* should acknowledge that when *h* is not possibly false in the serious sense directly pertinent to the conduct of deliberation and inquiry, it remains, nonetheless, logically possible that it is false (provided, of course, that *h* is not a logical truth).

Thus, in advancing the theses of epistemological and categorical infallibilism, I do not mean to counsel absurdity or dogmatism. On the other hand, if knowledge is to have the role as a resource in inquiry and deliberation which, in my opinion, it should have—namely, as a standard for serious possibility—it is of the first importance to understand that some elements of so-called fallibilist epistemologies require modification. My advocacy of epistemological and categorical infallibilism is intended to call attention to some of the necessary modifications.

I lack a demonstration that the epistemological outlook which emerges is correct. Yet, a conception of the uses of knowledge both in scientific inquiry and in practical deliberation is promised which avoids introducing the chasm between theory and practice to which both Peirce and Popper seem committed. For this reason alone, the epistemological outlook deserves a further hearing.

1.7 Truth as a Value in Inquiry

Peirce believed that advocacy of infallibilism places road-blocks in the path of inquiry. Rejection of corrigibilism creates obstacles. Categorical infallibilism does not.

Categorical infallibilism is no obstacle to the revision of a corpus of knowledge. Agents do and are capable of changing their minds. Moreover, they can justify such revisions provided the objectives they seek to promote in making revisions are of the right sort.

Peirce (and Popper too) maintained, however, that truth is a desideratum of scientific inquiry. Scientific methods of fixing beliefs are superior to others because a community of inquirers using such methods would, in the long run, converge on the true complete story of the world.⁵

For Peirce, the true complete story is true in the sense that it is the system of assumptions to which the community of inquirers using scientific methods would converge in the long

run. We need not endorse the conception of truth in this account in order to entertain as a serious view of the aims of inquiry the thesis that the ultimate goal is to attain a body of knowledge furnishing a true complete story. Popper endorses a Tarskian conception of truth. Yet he regards convergence on the truth as the ultimate aim of inquiry. He does not claim that we can be certain that scientific methods as he envisages them will succeed in attaining that goal. Yet he does claim that scientific methods are the best available means for attempting to attain that objective.⁶

In chapter 3, I shall attempt to explain why advocates of either Peirce's or Popper's version of the view that the ultimate aim of science is convergence on the true complete story would argue that corrigibilism requires rejection of categorical infallibilism and, hence, epistemological infallibilism together with the thesis that knowledge is the standard for serious possibility. If this analysis is on the right track, the heart of the controversy between fallibilism of the Peirce-Popper variety and corrigibilist infallibilism of the sort I am defending turns on a controversy concerning the aims of inquiry. I reject the view that convergence on the true complete story of the world is the ultimate aim of efforts to improve knowledge.

Nonetheless, avoidance of error is an important desideratum of the proximate aims of specific inquiries concerned with specific revisions of standards for serious possibility or corpora of knowledge.

Insofar as X is worried about importing error into his corpus by adding h to it, he should have no worries if h is already in his corpus. Nor should he be concerned to remove h from his corpus in order to avoid error. From his point of view, there is no serious possibility that h is false.

There may, indeed, be good reasons for removing h from his corpus. Such reasons will be examined in chapter 3. For the present, it is sufficient to note that avoidance of error cannot furnish a warrant for removing h .

X would have excellent reason for avoiding the importation of h into his corpus in case $\sim h$ were already in his corpus. From X 's initial point of view, $\sim h$ is infallibly and certainly true and h infallibly and certainly false. If X were concerned to avoid error, he would attempt to avoid importing h into his corpus.

Suppose neither h nor $\sim h$ is in X 's corpus at t . From X 's initial point of view, importing one or the other hypothesis into the corpus runs a risk of error. In considering what conclusions to draw, X should take these risks seriously and should be willing to incur them only if there is sufficient compensation promised from other cognitive goals.

Thus, X should seek to avoid error in revising his corpus at t relative to a definition of "true in L " (expressible in L_1) which secures the truth in L of all items in X 's corpus of knowledge at t expressible in L .

To be sure, "true in L " could be defined relative to UK and held fixed. As I describe the situation, that truth definition remains a part of the definition to which X is committed at t . As X revises his corpus of knowledge, he also revises his truth definition for L .

I am not, however, interested in verbal hocus-pocus. I am attempting to explain the sense in which avoidance of error is an invariant feature of what should be the proximate goals of efforts to revise corpora of knowledge by adding new information to them. I am advancing a thesis, alternative to the Peirce-Popper view, about what ought to be the aims of inquiry.

There is no single objective which all special inquiries seek to promote. One reason (by no means the only one) is that as knowledge changes, the respects in which avoidance of error is a desideratum in inquiry also change. Even if it were conceded to Peirce and Popper that, from X 's point of view, the ultimate aim of inquiry is to obtain a true, complete story of the world, it would still be the case that after X has revised his corpus of knowledge, his conception of what constitutes a true complete story alters as well.

It may be objected that " h is true" has been equated with " h is known or assumed by X to be true." That is not so. From X 's point of view at t , everything X knows at t is true. It is not the case that from X 's point of view, everything that is true is known by X at t .

Someone might complain that truth has been relativised to persons and times. I have, indeed, assumed that X 's knowledge is knowledge for X at t and have insisted that X 's theory of truth is a system of truth conditions adopted by X at t . In

this sense and in this sense alone is the charge of relativity correct.

However, the truth conditions adopted by *X* at *t* supply a definition of "*h* is true in *L*" and not of "*h* is true in *L* for *X* at *t*." Truth in this work remains as atemporal and objective as it is in any other Tarskilike conception.

Thus, when *X* is interested in modifying his knowledge while at the same time avoiding error, the revisions he makes should not be regarded by him or anyone else as involving his "making true" or "making false" various sentences. In modifying his corpus of knowledge, *X* is changing his system of assumptions as to what sentences are true and thereby changing his commitments as to the conditions under which sentences in *L* are true.

Have we now so far lowered our sights as to settle for a relativistic doctrine of truth rating the statements of each theory as true for that theory, and brooking no higher criticism? Not so. The saving consideration is that we continue to take seriously our own particular aggregate science, our own particular world theory or loose total fabric of quasi-theories, whatever it may be. Unlike Descartes, we own and use our beliefs of the moment, even in the midst of philosophizing, until by what is vaguely called scientific method we change them here and there for the better. Within our own total evolving doctrine, we can judge truth as earnestly and absolutely as can be, subject to correction, but that goes without saying.⁷

My only reservation concerning this admirable passage from Quine is that earnest judgment of truth is something more than earnest affirmation or earnest conviction. Our earnest judgments of truth furnish us with the characterization of error in the sense in which error is a desideratum in inquiries aimed at correcting our own total evolving doctrine.

My rejection of the Peirce-Popper view of the aims of inquiry runs much deeper than the comments just made concerning the role of *X*'s corpus of knowledge in specifying truth conditions indicates. To seek a true complete story of the world as the ultimate aim of inquiry implies that all revisions of knowledge should be appraised in terms of their long-run tendencies to yield a maximally informative and error-free corpus of knowledge. It is open to an advocate of the Peirce-Popper position to concede that *X*'s conception of the ultimate

aim of inquiry changes with revisions of his knowledge and yet insist that seeking a true complete story is an invariant feature of *X*'s ultimate aim at each stage of inquiry.

I mean to deny this view. Avoidance of error is an invariant feature of the proximate goals of inquiries concerned with the improvement of knowledge. That is to say, in contemplating revisions of knowledge, *X* should desire to avoid error in making the revision immediately contemplated. Speculation as to the prospects of that revision leading to the avoidance or importation of error in further revisions which might take place later on are not to be taken into account. On my view, *X* should not be concerned as to whether the revisions he is about to make will or will not import error in the long run. He should be concerned with whether the revision he is about to make will immediately import error or not. But once he has made the revision and his conception of error has changed, he will no longer be concerned with avoidance of error according to his initial view. His new conception of error will take over and guide revisions at the next stage.

Once this view of the role of avoidance of error as a desideratum in the proximate aims of inquiry is adopted, there is, in my opinion, no serious obstacle to endorsing the thesis of the use of knowledge as a standard for serious possibility (both in scientific inquiry and in practical deliberation), the corollary of epistemological infallibilism and, in addition, the thesis of corrigibilism. I shall attempt to show this to be so, at least in outline, in the two chapters which follow.

The Peirce-Popper view, according to which convergence on the truth is the ultimate aim of inquiry, precludes, as I shall argue, consistent advocacy of both corrigibilism and epistemological infallibilism. As I have stated, rejection of epistemological infallibilism leads both Peirce and Popper to what is, in my view, an untenable separation between theory and practice and to an obscure (I suspect nonexistent) conception of the uses of knowledge in inquiry and deliberation.

In addition, respect for truth becomes a value in inquiry placed on a remote pedestal to be worshipped from afar. Avoidance and risk of error seem to play no role for either philosopher in the immediate concerns of inquiry. By way of contrast, I seek to develop an account of the improvement of

knowledge which sees respect for truth as vitally relevant to the immediate objectives of science.

Thus, for Popper, the best method for converging on the truth (i.e., an error-free maximally consistent story) is to devise testworthy hypotheses, test them, and if they survive such testing, test them again. Nothing in Popper's criteria for testworthiness indicate that, save for the minimal requirement of consistency, avoidance of error is a serious consideration in the appraisal of hypotheses as testworthy. In specific inquiries, a scientist's immediate aim should not, according to Popperian doctrine, take avoidance of error into account at all. Truth enters into the picture only insofar as in some way (which I do not understand) the persistent and serious testing of highly testworthy hypotheses leads to ultimate convergence to the truth. In the name of objectivity, truth is to be worshipped from afar, but it should not intrude into the real business of science.

Popper himself has acknowledged and, indeed, insisted that testing hypotheses requires "background knowledge."⁸ Prima facie, the need for background knowledge derives from the importance of restricting the space of serious possibilities more tightly than taking all logical possibilities would do. However, precisely for this reason, it ought to be the central concern of an epistemology focused on the "growth of knowledge," as Popper's allegedly is, to devise criteria for modifying such knowledge. Otherwise, it may appear that revisions of background knowledge are not subject to critical control at all and surely not to criticism which takes avoidance of error seriously.[†]

Popper has contributed nothing to the question of the revision of background knowledge. He has, instead, emphasized the testworthiness of hypotheses on the assumption that we devise hypotheses in order to test them. But surely we often test hypotheses in order to modify such background knowledge so that it can be used in subsequent inquiries in science and to guide the conduct of daily life. And presumably the point of such tests is to keep the risk of error incurred by the subsequent revisions at an acceptably low level.

[†] I suspect that for Popper what constitutes "background knowledge" depends on the "situation" in the sense which certainty allegedly does according to the passage cited from *Objective Knowledge* in the footnote on p. 4.

Popper's view concerning the ultimate aims of inquiry prevent him from conceding this. Although he flirts with inconsistency by acknowledging the role of background knowledge in testing hypotheses, he does not quite concede that background knowledge serves as a standard for serious possibility. Hence, he avoids epistemological infallibilism in order to endorse both corrigibilism and categorical fallibilism. That is why Popper has failed to take the revision of background knowledge seriously. There is no obvious way in which he could focus on this matter without threatening his commitment to both categorical and epistemological fallibilism or his conception of the ultimate aims of inquiry.

1.8 Four Types of Revision

Given a suitably regimented language L and the hierarchy of metalanguages generated by it, all changes in X 's knowledge expressible in L can be represented by changes in deductively closed sets of sentences in L all of which contain UK . Four kinds of changes are recognizable.

- (i) *Expansion*: A shift is made from K_1 to K_2 containing K_1 obtained by adding a sentence e (or a set of sentences) to K_1 and forming the deductive closure.
- (ii) *Contraction*: A shift is made from K_1 to K_2 , where K_1 is an expansion of K_2 .
- (iii) *Replacement*: A shift is made from a consistent K_1 containing h to a consistent K_2 containing $\sim h$.
- (iv) *Residual Shift*: One that belongs to none of the other three categories.

Expansion occurs when X adds information to his corpus via observation or testimony furnished by individuals X assumes to be truthful and competent. It also occurs when X concludes that his current corpus (including "background knowledge" and data obtained via observation and testimony) furnishes sufficient warrant for adding a law, theory, or prediction to his corpus and he converts an erstwhile hypothesis into an assumption he is committed to using as part of his standard of serious possibility in subsequent inquiry.

Contraction occurs when X concludes that some item in his corpus ought to be subjected to critical scrutiny and test and

should, therefore, cease being an assumption or evidence and should become, instead, a hypothesis.

Replacements take place when X is converted from a commitment to one theory to a commitment to another conflicting with the first. This kind of shift has received the lion's share of attention in recent years from those authors who have been concerned with the "revolutionary" character of important changes in scientific knowledge.

Residual shifts have received little attention in the literature, and are mentioned here for the sake of completeness.

It is customary to represent changes such as these by sentences which are added or removed from a corpus. However, according to the view I am presenting, revisions of knowledge ought to be represented, strictly speaking, by changes of sets of sentences.

For example, "inductive inference" is often taken to be a mode of expansion whereby a generalization is added to a corpus containing, among other things, a set of reports of so-called confirming instances of the generalization. However, the expansion requires that the deductive consequences of the generalization should be added in addition to the generalization itself. The "data" do not "confirm" the generalization but confirm the addition of a set of sentences. We can, to be sure, represent what is confirmed by a single sentence—to wit, a sentence which, when added to the corpus, generates via deductive closure the entire set of new sentences to be added. Such a sentence is a strongest sentence whose addition to the corpus is warranted and, in that sense confirmed, by the data and other information already in the corpus.

Imagine that X had in $K_{x,t}$ a set of statements of the form " a is neither a raven nor black" and no sentences of the form " a is a raven." Let us concede, for the sake of the argument, that X should expand his corpus by adding "everything is a nonblack nonraven" along with its deductive consequences. Perhaps this is questionable. What is not questionable, in my opinion, is that if the data as described warrant adding "everything is a nonblack nonraven," they also warrant adding "all ravens are black."

Under typical conditions, it would be puzzling if "data" consisting exclusively of nonblack nonravens warranted adding "All ravens are black" but failed to support adding "every-

thing is a nonblack nonraven." That is to say, it is implausible that under normal circumstances such data warrant adding "all ravens are black" as a strongest added statement to a corpus although it is quite plausible that its addition to the corpus is warranted if the addition of "everything is a nonblack nonraven" is warranted. In my opinion, what has gone wrong in discussions of the so-called paradoxes of confirmation has been that the "conclusion" to be supported by data has always been taken to be a sentence or proposition rather than a set of sentences or propositions added to a corpus to make a deductively closed set.†

All feasible bodies of knowledge are expansions of or identical to UK . UK is itself noncontractible. This amounts to saying that all items in UK are incorrigible items of every feasible corpus expressible in L . Hence all items in UK are certainly and infallibly true for all agents at all times (if they are rational). However, if $K_{x,t}$ contains items not in UK , then from X 's point of view at t , there are sentences not in UK which are quite as certainly and infallibly true as the assumptions in UK . Incorrigibility implies infallibility. The converse does not hold. Hence, from X 's point of view at t , the theories, laws, statistical claims and observation reports he accepts in his corpus are as infallibly and certainly true as truths of logic, set theory, and mathematics. Certainty and infallibility are one thing. Corrigibility is another.

The feasible or potential corpora in L can be partially ordered by the set-inclusion relation. Moreover, it is plausible to require that for X , K should be at least as informative as K' if K' is a subset of K . K rules out more logical possibilities than K' . The least informative corpus is UK itself. The most informative one is the contradictory corpus.

Nothing prevents a contradictory corpus from being feasible. There is some value in being able to consider both changes in knowledge which end up in contradictory corpora and changes which move away from contradiction.

To allow X to consider a contradictory corpus to be feasible does not imply that if he should detect inconsistency in his

† For further elaboration and qualification of this approach to the so-called paradoxes of confirmation, see my "A Paradox for the Birds" in *Essays in Memory of Imre Lakatos* (edited by R. S. Cohen, P. K. Feyerabend, and M. W. Wartofsky. Dordrecht: Reidel, 1976, pp. 371-378).

corpus he should rest content. When X 's corpus is inconsistent, it breaks down as a standard of serious possibility. It furnishes a truth definition which is unsuitable for characterizing the aim of avoiding error. It is useless as a resource for inquiry and deliberation.

Moreover, should X contemplate expanding into an inconsistent corpus from one which is consistent, he would be well advised not to do so; for, from his initial point of view, such a move would deliberately import error into his corpus. For this reason, respect for truth would argue against expanding into contradiction. (This observation will be somewhat qualified in the next chapter.)

We should distinguish between commitments imposed on X by his having a corpus of knowledge at a given time by principles of epistemic logic specifying "equilibrium conditions" and those commitments to improve his corpus generated by X 's epistemic or cognitive aims and values. I propose to regard the desirability of moving away from contradiction and avoiding moving into contradiction as derived from commitments of the second kind. Consistency, on this view, is not a condition of rational equilibrium.

1.9 True Justified Belief

Objections may be raised against the cavalier downgrading of the classical formula that knowledge is true, justified belief. It may be claimed that some variant of the formula is needed to accommodate the widely acknowledged distinctions between knowledge and true but unjustified belief or between knowledge and justified false belief.

I do, indeed, reject the value of such distinctions when made by X at t concerning his beliefs at t . From X 's point of view at t , there is no difference between what he fully believes at t and what he knows at t . From his point of view at t , if he fully believes h at t , the falsity of h is not a serious possibility for him at t . Hence, according to X , h is true. There is no need for justification. How could there be, given that the falsity of h is not a serious possibility? Only if one insists that in order to claim knowledge that h , X must also show that he admitted h into his corpus legitimately, could an opening be made for a distinction between his knowing that h and his believing that h . But it is precisely this sort of pedigree epistemology that I mean to reject.

Nonetheless, X can consistently distinguish between what Y knows and what Y believes or between what X himself did know or will know and what he did or will believe.

More importantly, even from X 's point of view, X can ask whether he is justified in revising his current corpus in one way or another and can, from his current point of view, recognize that he may be justified and yet import error into his corpus.

Thus, in contemplating an expansion of his corpus by adding h and its deductive consequences, X can recognize as a serious possibility that in doing so he will admit error into his corpus. Nonetheless, he may, as I shall argue later, be justified in doing so. Hence, prior to expansion, X can recognize a significant distinction between "coming to know that h " when his expansion is justified and error free, coming to have justified but erroneous belief that h and coming to have error-free but unjustified belief that h .

Notice, however, that once he has expanded by adding h to his corpus, from his new point of view, h is certainly and infallibly true. Hence, as long as h remains in his corpus, from his point of view, X knows that h even though his initial expansion happened to be illegitimate. To be sure, if he comes to recognize somehow that he committed an error in expansion or that the expansion was unjustified, this may prove relevant to a decision as to whether to contract. However, it need not be automatically decisive in favor of contraction; for, it should be kept in mind, that as long as h is in X 's corpus, he considers h to be infallibly and incorrigibly true. In chapter 3, I shall elaborate on the topic of contraction at greater length. The point I wish to emphasize now is that the distinction between knowledge, true belief, and justified belief does have some relevance when X reflects on his own beliefs at other times or on Y 's beliefs. It also has importance when revisions of belief are considered. But it has no relevance from X 's point of view for the purpose of making discriminations within his corpus at t .

Let X 's corpus at t_1 be K_1 . He adds h to that corpus. Suppose he is justified in doing so. In virtue of deductive closure, his corpus K_2 at t_2 contains not only h but $h \vee g$. At t_1 when X elected to expand, he did not know whether in expanding he would avoid error. Hence, even though he had

concluded that in expanding, he would come to justifiable belief that h and that $h \vee g$, he could not claim that he had come to know that h .

Suppose that Y (at t_2) has $\sim h$ in his corpus. He concedes that X was justified relative to K_1 in expanding the way he did. Moreover, Y has g in his corpus so that, from Y 's point of view, $h \vee g$ is true though h is false.

Edmund Gettier asked: Does X know $h \vee g$? From X 's point of view after expansion, it is obvious that he does. What about from Y 's point of view?

There are two distinct questions: Did X come to know $h \vee g$? In my view, the answer is negative. In expanding from K_1 to K_2 , error is imported if a single false proposition is added. If so, X has failed to come to know that $h \vee g$ from Y 's point of view because in expanding, X has added false h to his corpus.

On the other hand, Y can consider the matter from an entirely different angle. He may ask what X should do in order to remove error from his corpus. Clearly, Y would recommend that X remove h from his corpus. Would he recommend that X remove $h \vee g$? Given Y 's opinions, it would be clearly foolish for X to do so. He would be surrendering true information. X should not cease to believe $h \vee g$ either from X 's point of view or Y 's. From both points of view, he is justified in retaining it. Moreover, from both points of view, $h \vee g$ is true. In what sense, therefore, should Y claim that X does not know $h \vee g$? Only in the sense that he failed to come to know it. But that observation is irrelevant, once we have ceased worrying about pedigree, in determining whether he should give $h \vee g$ up. In this sense, Y can in good conscience concede that X knows that $h \vee g$.

I do not offer this account as an analysis of the presystematic usage of the verb "to know." There are many terms, and "know" is one of them, where presystematic usage and our "intuitions" about them are heavily burdened with the commitments of philosophical ideology; so that consultation with intuition rarely avoids begging the issues under dispute. I put this account of the Gettier problem forward only to illustrate how puzzles such as this appear from the perspective of the epistemology I have been outlining here.

1.10 Free Speech

Advocacy of epistemological and categorical infallibilism appears to undercut some liberal arguments in favor of free speech and open criticism. If from X 's point of view at t (whether X is a person or a social group) there is no serious possibility that h is false, then with respect to any decision to be taken to which the truth value of h is relevant X might regard dissent from his (or its) views concerning the truth value of h as an impediment to making the right decision. Advocating infallibilism seems to furnish part of a grounding for circumscribing free speech.

Even if it were true, this complaint would not suffice to make the case for endorsing categorical fallibilism; for even if X were obliged to be uncertain about the truth of every extra-logical statement, he would be guided in his conduct by his discriminations between hypotheses with respect to degrees of uncertainty. If anyone dissents from his appraisals and, hence, his judgments concerning the policies which should be adopted, from X 's point of view the dissent would remain an impediment to rational choice. If a case for circumscribing free speech and criticism can be sustained by categorical infallibilism, it could be sustained by invoking categorical fallibilism as well.

The alleged implication of categorical infallibilism is invalid. Whether tolerating or promoting free speech and dissent is undesirable from X 's point of view should depend on the usefulness, from X 's point of view, of institutional arrangements protecting free speech and criticism in promoting improvements in X 's knowledge and, through such improvements, in furnishing a better basis for policy making. To make a case for free speech and inquiry on these grounds, X must be committed to the view that his knowledge is open to improvement. Corrigibilism is a necessary assumption. But corrigibilism is not to be confused with categorical fallibilism as such exponents of the liberal view as Popper have done.

In any case, corrigibilism alone, though not an impediment to defending free speech, is not a sufficient basis for endorsing it. A case needs to be made that the improvement of knowledge is, in point of fact, promoted by tolerance and encouragement of dissent; and this, for some, may remain an open question.

It does seem clear that if in science or government reliable

information is to be obtained, sufficient freedom must be granted to scientific investigators or government agents (e.g., spies and diplomats) so that they can furnish the information requested. Governments are prone, however, to circumscribe such freedom so that agents can perform the functions demanded by government but cannot exercise their freedom in other domains. Policies of the Soviet Union towards scientists and artists furnish a familiar example of governments confronting this problem; but the efforts on the part of the United States government to restrict the dissemination of "classified information" or the capacity of government employees to engage in public criticism of some aspects of government policy illustrate the same point.

Of course, from the point of view of private citizens, the problem of obtaining useful information for the making of public policy as well as for taking private decisions will be seen in a different light. The question of what is an optimal allocation of freedoms for the purpose of obtaining reliable information will be considered in different ways depending on the sorts of information being sought.

Given that even X (whether X is the government, a scientific community, or a private citizen) can recognize the existence of conflicts between his needs on this score and the needs of other agencies, the question of determining how to allocate freedom of speech and criticism to individuals and groups in society can be recognized by X to be a moral and political problem of optimal distribution.

Here I do not intend to go into this topic any further. I am concerned only to show (a) that liberal conclusions concerning the importance of promoting free speech and criticism are neither entailed by nor undermined by advocacy of categorical infallibilism, (b) that commitment to corrigibilism appears more critical to the liberal point of view, and (c) that even such a commitment to corrigibilism will not, without appeal to other moral and political assumptions concerning just distributions of freedoms and without careful scrutiny of how various institutional frameworks do and do not facilitate the improvement of knowledge, furnish a sufficient grounding for the liberal conclusion.

I am sympathetic with this conclusion; but it seems to me that efforts on the part of authors like Mill and Popper to

buttress their liberal convictions by appeal to their epistemologies (and vice versa) do not succeed.

1.11 Is Revision Possible?

I have claimed that corrigibilism is compatible with both epistemic and categorical infallibilism provided that seeking the truth is not taken to be the ultimate aim of inquiry though avoidance of error is regarded as a proximate desideratum of particular inquiries. To support this claim we must consider how X can reasonably modify his corpus of knowledge given that his initial corpus is, as far as he is concerned, infallible and given that among his epistemic goals in modifying his knowledge is a concern to avoid error.

This question can now be seen to have four parts:

- (a) How is expansion possible?
- (b) How is contraction possible?
- (c) How is replacement possible?
- (d) How is a residual shift possible?

These questions will be considered in the two chapters which follow.

2.1 New, Error- Free Informa- tion

The idea that scientific progress is marked by a steady accumulation of knowledge is an object of contemporary scorn. Thanks to the efforts of T. Kuhn and P. K. Feyerabend it is widely held that in science theories are often overthrown and replaced by other theories. Yet, Kuhn insists that individuals and scientific communities often modify their knowledge in a manner conforming to a cumulative model. *Prima facie*, knowledge is modified by expansion as well as by contraction and replacement.

Why should an agent expand his corpus? The short answer is: To obtain new error-free information.

The demand for new information will vary from context to context, and will often be motivated by strikingly different considerations on different occasions. Not only will the subject matters of different inquiries vary; the inquiries themselves will be motivated by the desire either to resolve some practical question, to contribute to systematic understanding, or both.

Whatever the demand for information might be, and whatever the reasons for that demand, it should be tempered by a concern to avoid injecting error into the corpus through expansion. In nontrivial expansion, information not contained in $K_{N,t}$ is added to it. From X 's point of view at t , that information is possibly false. A scientifically responsible X should be concerned to avoid error. The proximate aim of his efforts to expand his corpus should be to obtain error-free information.

How could expansion be justified for the sake of obtaining new error-free information? Prior to expansion, X regards all items in his corpus as infallibly true. Hypotheses not in his corpus but consistent with it are fallible—i.e., possibly false. If X should add h to his corpus, he would be adding a fallible hypothesis—counter to the claim that what is in his corpus is infallible.

This objection presupposes that possibility or infallibility belongs to sentences or propositions independently of the agent's corpus of knowledge. Even if this were so for logical possibility and other sorts of possibility considered in the literature, it is false in the case of serious possibility of the sort relevant to the conduct of inquiry and deliberation that is being discussed here. If X changes his corpus, he changes his standard for evaluating serious possibility and, hence, what is infallible for him. From X 's point of view prior to expansion, adding h to his corpus is, indeed, a possible source of error. Accepting h is fallible *ex ante*. It is infallible according to X *ex post*. Once X has adopted h as an item in his corpus and shifted his point of view, h becomes infallible for the purpose of subsequent inquiry and deliberation.

There is, however, a more serious objection to expansion. From X 's point of view at t , h is possibly false. Adding h to his corpus entails a risk of error. Refusing to expand at all incurs no risk. Surely the desirability of avoiding error favors the latter option over expansion.

Were the proximate aim of an effort to expand simply the avoidance of error, this point would be decisive. Yet X should be concerned not merely to avoid error but to acquire new information. The promise of obtaining new information may sometimes (though not always) compensate X , from his initial point of view, for the risk to be incurred.

2.2 Routine and Inferential Expansion

Expansions are inferential or routine. The former are a species of deliberate decision making. The latter are instances of routine or programmed behavior.

Tom is driving to destination B and has reached a fork in the road. He does not know whether to turn left or right. He decides quite deliberately that he will turn left if a coin he intends to toss lands heads up and right if the coin lands tails up. The coin lands tails up and Tom turns right.

Tom does not deliberately choose to turn right. He does deliberately choose to let the outcome of the toss determine the direction in which he turns. With respect to the issue as to how he turns, Tom has acted in a routine or programmed fashion. With respect to whether he shall follow a program and if so which program he will follow, Tom has acted deliberately.

Dick is driving to the same destination and has reached the same fork in the road. There is a public telephone at the side of the road. Dick calls the AAA to ask for advice having decided in advance that he will do whatever they tell him. They tell him to turn right and so he does.

Dick does not deliberately choose to turn right any more than Tom does. He does deliberately choose to let the outcome of his query to the AAA determine the direction in which he turns. He too has acted in a routine or programmed fashion with respect to this issue. But he has chosen a different program than Tom and has chosen it deliberately.

Harry is also travelling to *B* and reaches the same fork in the road. He, like Dick, makes the telephone call to the AAA and receives the same instruction after having decided in advance to follow the instructions whatever they might be. Upon returning to his car, Harry notices a road sign obscured by the trees which states that the road to *B* is the road to the left—counter to the AAA instructions. Harry ignores the sign because he has committed himself already to following the AAA instructions. He too chooses to let the outcome of some process (interrogating the AAA) determine the direction he travels. He does, however, deliberately choose the routine.

Had Harry, upon seeing the road sign, weighed the reliability of the AAA against the reliability of the signs and, perhaps, sought more data in order to resolve the conflict before turning in one direction or the other, Harry would have deliberately chosen the direction in which he was to travel. In deliberate decision making, the agent identifies the options available to him, his goals, and the available relevant evidence concerning the admissibility of the options for the purpose of realizing these goals and values. The option chosen is determined relative to these beliefs and values according to principles of rational choice.

Such cases of deliberate choice contrast with situations where action is the “output” of a process in which the agent responds to “inputs” according to some program. The program may be innate, acquired by conditioning, or may itself be chosen as the outcome of deliberation. Following such a program is letting the action to be performed be settled as the outcome of some process which is often stochastic. The agent,

in effect, treats himself as part of a “chance setup”[†] being subjected to some sort of trial and lets the act performed depend on the outcome of the trial.

To be sure, deliberate decision making itself is a process whose output is an act and whose inputs are other factors such as beliefs and values. Conceding the point does not undermine the distinction between deliberate and routine decision making. The “program” for deliberate decision making ought to be whatever general principles qualify as principles of rational decision making determining which options are admissible given specific information concerning the options the agent recognizes to be feasible for him, given his values and goals and given his beliefs. To characterize deliberate decision making is to furnish such principles of rational choice and to identify (within the limits of precision that are feasible) the domain of applicability of such principles. Much of this book is devoted to supplying materials which should help clarify these matters at least in part. Given such an account we would be in a better position to give a precise characterization of the distinction between deliberate and routine decision making. Lacking such an account at present, I shall take for granted that the distinction is legitimate and important.

Routine expansions are instances of routine decision making. Thus, *X* may follow the practice of consulting *Y* on matters concerning some domain. In interrogating *Y*, *X* is conducting a trial on a chance setup whose output is *Y*'s report. *X* adopts a program utilizing that report as an input into another program whose output is *X*'s adding *Y*'s report to his corpus. The combined program amounts to employing a certain kind of trial (interrogating *Y*) where the possible outcomes are modifications of *X*'s corpus by expansion.

Another important kind of routine expansion involves *X*'s making observation reports in response to sensory stimulation and adding the reports made (or corrections of them) to his corpus.[‡]

[†] The term “chance set-up” was introduced by I. Hacking in his *Logic of Statistical Inference* (Cambridge: Cambridge University Press, 1965, ch. II). See also section 11.20 of the present book.

[‡] Making the observation report that *h* is not to be confused with adding *h* to one's corpus via expansion. In making observations, *X* typically lets the outcome of sensory stimulation (i.e., the making of the observation report

Routine expansion via observation and routine expansion via the testimony of others involve X 's adopting, prior to his conducting a "trial" and his ascertaining the result, some program which commits him to expand in a manner which depends on the observation report he makes in the case of routine expansion via observation or on the testimony of Y in case the program concerns expansion through consulting others. He is committed by adopting the program to carrying it through.

In contrast to the foregoing, inferential expansion is a species of deliberate decision making. In deliberate decision making, X identifies a set of feasible options and undertakes to determine which of them is admissible given all of the information available to him prior to making his choice and given his goals and values. Of course, inferential expansion is not deliberate decision making aimed at realizing ethical, economic, political, or other such practical goals. As in routine expansion, the aims are cognitive. The options are potential expansion strategies which qualify as potential answers to the question under investigation, and the aim is to gratify the de-

that h) decide what he will add to his corpus. Often he does so as a matter of habit. But in carefully designed experimental situations, the circumstances under which he will let the outcome of sensory stimulation render a verdict are circumscribed. If X knows that observations made on a foggy night have a great chance of yielding false reports, he will refuse to add foggy-night sentences, even though they have been reported true, into his corpus.

The sentence added to X 's corpus via observation routine is not inferred from what is already in his corpus (including knowledge of the error probability of the routine) and information concerning the response made to sensory stimulation. Rather, X is committed prior to observation to the practice or routine of letting the application of the routine legislate what he will add. Being so committed, the actual implementation of the routine is not an inference at all.

The structure of expansion via observation resembles the structure of statistical decision making as construed by advocates of the Neyman-Pearson approach to statistics. According to that view, an investigator plans beforehand which outcomes of an experiment will lead to rejection of a "null hypothesis" and which will not. The probability of error (and other "operating characteristics," such as the power of the test) are determined relative to the knowledge available prior to the experiment, and the plan is evaluated on the basis of this information. This procedure stands in contrast to approaches according to which one decides whether one should reject the null hypothesis relative to a body of knowledge including information about the outcome of the experiment. This contrast is discussed in somewhat greater detail in chapter 17.

mand for information occasioned by the question while at the same time avoiding error.

Thus, an investigator seeking the true value of some parameter might regard any point estimate as a potential answer. He will also consider weaker claims asserting that the true value falls in some interval or in one of a set of intervals. Or the investigator might be seeking a theoretical basis for systematizing some subject matter and might be entertaining the alternatives of adding theory T_1 to his corpus, adding T_2 , rejecting both theories (i.e., adding the residual hypothesis R which asserts that both are false), suspending judgment between T_1 and T_2 , between T_1 and R , between T_2 and R , and between all three alternatives.

Given the demands of his question and the potential answers he has identified, X should attempt to justify adopting one potential answer rather than another utilizing all the knowledge already available in his corpus.

By way of contrast, in routine expansion X does not choose one from a list of potential answers. He may not be seeking an answer to some specific question (although he could be) and he need not have identified a list of potential answers. X has been conditioned or committed to a program for selecting expansion strategies depending on the outcome of a trial of some kind but not depending on his finding out what that outcome is and combining it with the rest of his knowledge to decide upon an optimum expansion strategy.

The aim of both routine and inferential expansion is the acquisition of new error-free information. However, the way concern to avoid error and to obtain new information trade off is different in the two kinds of expansion.

Thus, if, in response to sensory stimulation, X makes an observation report inconsistent with what is already in his corpus and, hence, certainly false from the point of view he adopted prior to carrying out the observation routine, X is committed to following through and admitting the report to his corpus. This is so even though carrying through converts his corpus to an inconsistent one.

Following through in this way does not trample on the desideratum of avoiding error. When X committed himself to the routine, he might have assumed that the routine had a small chance of yielding an erroneous expansion (due to the

fallibility of *X*'s senses) but thought it justifiable to follow the routine because the chance of error was sufficiently small to be amply compensated for by the information promised by following the routine. Once he adopted the routine on this basis, the fact that, from his initial point of view, an erroneous report was made in the specific case under consideration means that in having elected to follow the routine beforehand he took a risk and lost. But it does not mean that *X* deliberately imported error into his corpus.

Similar remarks apply to routine expansion via appeal to the testimony of competent witnesses or experts. Suppose that *X* assumes *Y* to be competent and reliable in giving testimony on a certain topic. That is to say, *X* assumes that the chance of *Y* answering erroneously when questioned on matters pertaining to the topic is low. *X* might regard the information to be gained by importing *Y*'s testimony as a matter of course into his own corpus as worth the risk of error. Notice, however, that *Y* might sometimes testify in a manner contradicting items already in *X*'s corpus. Hence, *X*'s following the routine will lead to his expanding into contradiction. Nonetheless, *X* has not deliberately done this. The contradiction is obtained by *X* following a routine for expansion which he assumes is highly albeit not perfectly reliable. He has accepted a risk of error for the sake of new information and lost.

By way of contrast, in inferential expansion, *X* does not let some stochastic process decide for him what to add to his corpus. Instead, *X* compares rival expansion strategies with respect to both risk of error and informational benefits promised relative to all the information available to him prior to expansion and decides which strategy furnishes the best trade-off. Although the informational benefits promised will sometimes warrant risking error, no such benefits should warrant risking certain error. *X* should not deliberately import error into his corpus for certain. *X* should not deliberately expand into contradiction.

Thus, if instead of using *Y*'s testimony as input into a program to which *X* committed himself beforehand for expanding his corpus, *X* had first added the information that *Y* had made a report contradicting what is already in *X*'s corpus, *X* should not expand by adding *Y*'s report to his corpus expressible in *L*.

Neither the testimony of the senses nor of other witnesses added via an expansion routine are, once admitted to *X*'s corpus, distinguishable from theories, laws, statistical assumptions, predictions, or other singular hypotheses with respect to certainty or infallibility. From *X*'s point of view, all items in his corpus are equally certain and infallible. Nor do observation reports have a special status with respect to corrigibility. Finally, they are not distinguished by their content.

Observation reports are distinguished by the way they gain admission into a corpus. Like the testimony of competent or expert witnesses, they gain entry through the implementation of a routine. Unlike inferential or inductive expansion (I shall use these terms interchangeably in this book), routine expansion is capable of injecting contradiction into a corpus even when implemented in a way which respects the desideratum of avoiding error.

This trait is far from being a virtue. It is a defect of all modes of routine expansion that they can breed error and, indeed, lead to the contradiction of our most cherished theories. We put up with the defect because of the information to be gained by consulting the testimony of our senses and of other reliable witnesses. We ought not, however, make a virtue out of our necessity and maintain that our senses are the ultimate arbiter of what should and should not be in a corpus of knowledge. There is no such ultimate arbiter.

2.3 Abduction vs. Induction

In seeking new error-free information via expansion of his corpus, *X* tries to answer some question. Often, however, the question or demand for information is obscure and part of the task of inquiry is clarification of the question. The most important aspect of this effort will be identifying potential answers to the question under investigation.

The task is sometimes trivial. If *X* is estimating the average height of undergraduates at Columbia College in 1975, the range of potential answers is relatively easy to identify. Considerable genius may be required, however, to identify a theory worthy of consideration for the purpose of systematizing some subject matter.

Once, however, potential answers are proposed for the status of potential answers to the question under consideration, it is sometimes important to determine whether the hypotheses

proposed do indeed qualify as potential answers to the question under consideration. Also, even if they do to some extent respond to the demands for information occasioned by the question, it is also important to ascertain how well they gratify these demands when the question of truth value is set to one side.

In large measure, such appraisals of expansion strategies depend on the question being investigated and the demands for information occasioned by it. However, to some extent, these appraisals may be regulated by criteria applicable to members of a large class of demands for information. Thus, certain types of explanation-seeking question may be such that potential answers should meet certain conditions in order to gratify the demands for information to a high degree—i.e., in order to bear high informational value relative to the question being considered.

Criteria of this sort correspond, at least roughly, to what C. S. Peirce called principles of abduction.

At least on some occasions, Peirce does not regard abduction to be a mode of inference leading to the fixing of beliefs—i.e., he does not consider them to be a way of adding new items to a corpus of knowledge via inference to be used as premises in subsequent inquiry. The “conclusion” of an abductive “inference” is an evaluation of an expansion strategy as a potential answer to a given question.

It is to be remarked that, in pure abduction, it can never be justifiable to accept the hypothesis otherwise than as an interrogation. But as long as that condition is observed, no positive falsity is to be feared.¹

As Peirce notes, the “conclusion” of an abduction can entail no error; for such a conclusion is the mere entertaining of a hypothesis for further test, scrutiny, and inquiry in an effort to answer some demand for information. In inductive inference (i.e., inferential expansion), on the other hand, an erstwhile hypothesis which is already taken to be a potential answer to the question under investigation is added to the corpus of knowledge, its status as hypothesis is stripped from it, and it becomes a settled answer to the question. From *X*'s point of view prior to making this inductive inference, to do so does entail a risk of error.

In the context of abduction, the only factors which need to

be taken into account in appraising hypotheses are their informational virtues. For example, when the aim is to add theories to the corpus which furnish systematic explanations of phenomena in some domain, good potential answers are good potential explanations possessing such virtues as generality and simplicity. Neither the truth values nor the probabilities of potential answers are relevant to the evaluation of such potential answers as good potential explainers. Hence, there is no need to decide whether informationally attractive hypotheses are or are not true or are likely or unlikely to be true.

Hence, even if simplicity is an informational virtue, there is no need, in the context of abduction, to settle the question of whether nature is simple or not. To do so would, in any case, be dubious. The concept of simplicity is notoriously ambiguous and, even when it is disambiguated in ways which render it *prima facie* relevant in scientific inquiry, evaluations of hypotheses with respect to simplicity seem to depend on the question under investigation.

Even in situations where widely shared criteria for assessing simplicity can be invoked, as in some curve-fitting cases, it is not true that simpler hypotheses are more probable. Relative to data represented by two points, it is no more probable that the true hypothesis is a straight line passing through the two points than that it is some particular circle passing through those points. Fortunately, in the context of abduction we do not have to suppose otherwise. In that context, probability of truth and of error is irrelevant.

Inductive or inferential expansion is an entirely different affair. Informational desiderata relevant to the assessment of the quality of hypotheses as potential answers continue to be relevant when the concern is to choose one of the list of potential answers for addition to a corpus of knowledge. But the risks entailed by the various potential answers must be taken into account as well. Probability of error does matter.

Many authors seem to hold that the difference between abduction and induction is a matter of degree or that induction is a species of abduction. In effect, such authors claim that the same criteria apply to the evaluation of abductions and to the evaluation of inductions. They are, therefore, committed to one of the following two alternatives:

(i) Hypotheses which are informationally attractive—perhaps, because they are simple—are more probably true than those which are not.

(ii) Avoidance of error is not a desideratum of efforts at expansion.

In my opinion, both alternatives are untenable. Hence, on my view, abduction and induction do not merely vary in degree.

If I understand them correctly, W. V. O. Quine and J. Ullian disagree:

Calling a belief a hypothesis says nothing as to what the belief is about, how firmly it is held, or how well founded it is. Calling it a hypothesis suggests rather what sort of reason we have for adopting or entertaining it. A man adopts or entertains a hypothesis because it would explain, if it were true, some things he already believes. Its evidence is seen in its consequences.²

The reasons, according to Quine and Ullian, for “adopting or entertaining” a hypothesis are those which render the hypothesis a good *potential* explanation of what the investigator already believes. These reasons are grounds for identifying hypotheses as potential answers at the abductive phase—i.e., for “entertaining” them. But Quine and Ullian hold that these very same reasons or sorts of reasons are grounds for adopting hypotheses as premises or as beliefs. In any case, Quine and Ullian quite explicitly state that induction is but a species of “framing hypotheses.”³

Thus, Quine and Ullian do appear to hold that the same informational desiderata control without additional supplement the selection of one of a list of potential answers as a conclusion to be used as a resource in subsequent inquiry. If I understand them correctly, they are faced with the dilemma of assuming dubious claims such as the simplicity of nature or denying that avoidance of error is a desideratum of inquiry. In “On Simple Hypotheses of a Complex World,” Quine registers skepticism as to the claim that nature is simple.⁴ Except, perhaps, for avoiding contradiction, avoidance of error is no more a desideratum in induction than it is in abduction.

On the view I advocate, avoidance of error is relevant to

induction though not to abduction. This means that hypotheses which earn high marks when informational desiderata alone are taken into account may, nonetheless, fail to gain admission into X 's corpus because the risk entailed by such expansion is too great. In other situations, simplicity may be a sufficient inducement to incur the risk.

Thus, we do more than judge truth earnestly and seriously relative to our evolving doctrine. Our judgments of truth define for us our aims in revising our doctrine—or, at least, that aspect of our aims characterized by the desideratum that error be avoided. To be sure, our judgments of truth must be supplemented by judgments of probability of truth and of error. That is a topic which shall be discussed at great length in later chapters. But the relevance of probability of error to the evaluation of expansion strategies in induction is a manifestation of our concern to avoid error.

2.4 Potential Answers and Informational Value

The set of potential answers X has identified for the question under investigation may be represented by specifying a set U of hypotheses such that $K_{x,t}$ entails the truth of at least and at most one element of U and each element of U is consistent with $K_{x,t}$. Following terminology I have used elsewhere, the set U is an *ultimate partition*.⁵

A potential answer may be represented as a case of rejecting all and only members of a subset of U . That is to say, $K_{x,t}$ is expanded by adding the deductive consequences of the assertion that the true element of U is not in the subset of rejected elements. When no elements of U are rejected, the expansion strategy is a degenerate one where X does not modify $K_{x,t}$ at all but remains in full suspense between all alternatives in U . At the other extreme, X might reject all elements of U and deliberately expand into contradiction. If X rejects all elements of U but h_i , then X adds h_i to his corpus along with its deductive consequences. Such a potential answer is a strongest consistent potential answer. When more than one element of the ultimate partition survives rejection, X suspends judgment after expansion between the rival unrejected alternatives.

When the ultimate partition contains a finite number n of hypotheses, the number of potential answers, including the contradictory one, is 2^n . These may be partially ordered with

respect to informational value. If one potential answer involves rejection of all and only hypotheses in the ultimate partition in a proper subset of the set of hypotheses rejected according to another potential answer it is less valuable informationally than the second potential answer. Rejecting no elements of the ultimate partition is, therefore, the least valuable informationally of all the potential answers. Rejecting all of them (and, hence, contradicting oneself) is the most valuable informationally.

The ultimate partition can contain a countable or even a noncountable infinity of alternatives—e.g., in situations where the problem is to ascertain the true value of some real-valued parameter or n -tuple of real-valued parameters. In the case where the number of alternatives is countably infinite, the remarks made thus far concerning the finite case need not be modified. When the ultimate partition is noncountably infinite, rejecting elements of any subset might be considered a potential answer. I shall restrict attention to situations where only Lebesgue measurable subsets need be considered.⁶

The condition imposed on the assessment of informational value by the partial ordering cited above is, in effect, a principle of abductive logic. It is intended to be a principle that regulates the evaluation of information regardless of the demand for information under investigation. Similarly, the condition that the set of potential answers be generated by an ultimate partition can be construed as a principle of abductive logic.

However, neither of these principles exhaust what can be said about abduction. In particular, they do not imply that X 's set of potential answers is uniquely determined by his current corpus. In addition, the demand for information which occasions X 's question has to be taken into account. Such a demand controls in part what shall count as an element of an ultimate partition for the problem under investigation. The potential answers appropriate to demands for information about genetic linkage are obviously different from those appropriate to the study of business cycles.

Furthermore, the evaluation of the informational value of potential answers will also depend on the demands for information. If X is seeking a theory which will furnish systematic explanations in some domain and has identified two potentially

explanatory theories T_1 and T_2 , his ultimate partition may consist of these two rival (i.e., incompatible) theories and the residual hypothesis R which asserts that both T_1 and T_2 are false. Although R will be an element of the ultimate partition and adding it to the corpus is a potential answer, doing so is less satisfactory informationally than adding T_1 or adding T_2 . R is less satisfactory for the obvious reason that it fails to satisfy the demand for a theory which will systematize the domain under investigation. Furthermore, there may be important differences between T_1 and T_2 with respect to explanatory adequacy which may warrant assigning them different informational values. Assessments of this kind are part of the abductive task. To some extent, these assessments may be regulated by criteria which are applicable to a large class of problems. It may, perhaps, be possible to identify certain desiderata which determine explanatory power and simplicity relevant to the assessment of informational value in inquiries where the aim is to obtain explanations of some kind. It is doubtful, however, that such desiderata can be converted into criteria for the evaluation of informational value which render it irrelevant to consider the peculiarities of the particular demands for information motivating specific inquiries. Indeed, such restrictions on the assessment of informational value are likely to be very weak. Such assessment is, in my opinion, heavily context dependent.

My concern here is not, however, to explore the extent to which informational value is or is not context dependent. The question is an important one and I hope to explore it further elsewhere. I will only mention one proposal that I have already made as a restriction on all assessments of informational value, as it will be important in the subsequent discussion.

Under idealized circumstances, I propose that the evaluation of the informational value of potential answers can be represented in the case of an ultimate partition with a finite number of elements by a function which assigns to elements of the ultimate partition positive real values such that the sum of the values assigned elements of the ultimate partition equals 1. The informational value of a potential answer is the sum of the values assigned to the elements of the ultimate partition which are rejected according to that potential answer. The contradictory potential answer receives the maximum value of

1 and the potential answer according to which no element of the partition is rejected receives the informational value 0. The partial ordering with respect to informational value mentioned earlier is preserved by this type of evaluation.

The assignment of values to elements of the ultimate partition is represented by a function $M(h_i)$. This can obviously be extended to a probability function by assigning g , which is a disjunction of elements of the ultimate partition, an M -value equal to the sum of the M -values assigned elements of the disjunction (where each element in the disjunction is counted once and only once). The value $M(g)$ represents the informational value of *rejecting* g or, alternatively, of adding the disjunction of all elements of the partition which are not disjuncts in g to the corpus together with their deductive consequences.

This M -function is an *information-determining probability*. It represents X 's evaluation of the informational value of rival potential answers. The adequacy of the representation depends on the demands of the question under consideration for new information. It does not represent X 's assessment of the probability of error involved in adopting the potential answer. It is not what shall be later called an *expectation-determining probability* or *credal probability*. Its function in inquiry is quite different from that of credal probability (as it is also from *chance* or *statistical probability* and from *confirmational probability*, both of which will be discussed at length subsequently).[†]

I have no proof that informational value should be assessed in a manner satisfying the condition just cited. Indeed, in one important respect, I do not think that the condition should always be satisfied. Sometimes numerical representations of informational value are inappropriate. In section 8.6, I consider cases where X 's evaluations of informational value should be representable by convex sets of M -functions but not necessarily single M -functions. But the case where such evaluations are representable by a single M -function remains an important special one and much that can be said about it

[†] My view stands in opposition to opinions expressed by Y. Bar-Hillel, R. Carnap, R. Hilpinen, J. Hintikka, and K. R. Popper who use credal, confirmational, or "logical" probabilities as information-determining.

relatively simply applies with appropriate technical qualifications to the general case.

Sometimes the elements of the ultimate partition U are representable by indices which are n -tuples θ of real numbers corresponding to points in some region in an n -dimensional space. A potential answer involves concluding that the true value of θ falls in some subset of points in that region. The M -value of a potential answer specifying that the true value of θ falls in some subregion of the region in the n -dimensional space is characterizable as the value of $\int m(\theta)d\theta$ over points in that region, where the function $m(\theta)$ is a density function.⁷ The density function can be used in lieu of the M -function to represent the evaluations with respect to informational value.

When U is countably infinite or when we are dealing with continuous parameters ranging over all points in an n -dimensional space, new problems of a technical nature emerge. A common sort of problem is estimating the variance of a normally distributed random variable. Often the demands for information suggest that interval estimates of the value of the logarithm of the variance of equal length should be assigned equal M -value. This implies, however, that if the density is positive, the integral over the entire interval from $-\infty$ to $+\infty$ will diverge and if the density is 0 the total integral is 0. Similar problems arise when U is countably infinite and where each element of U is assigned equal M -value.

Difficulties such as these are generated by the peculiarities of infinity. They are important and should not be ignored; but they should not be raised as objections to the basic idea of representing informational value by means of information-determining probability measures. Parallel problems arise for expectation-determining probabilities and I shall discuss them in section 5.11. The suggestions made there can be adapted to the representation of information-determining probability.

Setting the reservations and qualifications just mentioned to one side, my proposal is that the results of abduction should be the specification of an ultimate partition and a set of potential answers generated thereby together with an information-determining M -function characterizing the informational values of the various potential answers. This proposal, it should be emphasized, imposes rather weak conditions on the results of abduction. It does not answer questions as to how an M -

function is to be selected and how the specific demands for information occasioned by an inquiry at hand control that selection. For the present I intend only to provide a relatively abstract representation of abduction to be used in an equally abstract characterization of criteria for evaluating potential expansion strategies, that in turn determine which of a set of potential answers can be adopted at the inductive stage.

2.5 Epistemic Utility

If X 's aim in inferential expansion were solely to obtain new information pertinent to the question under consideration, the evaluations of the potential answers generated by the M -function would provide a sufficient basis for deciding which potential answer to adopt. The informational value $M(g)$ of rejecting all elements of the ultimate partition which are disjuncts in g is the value of endorsing that expansion strategy relative to the aim of obtaining new information pertinent to the question without regard for the avoidance of error. This value accrues to the potential answer regardless of whether in rejecting g X does so erroneously (i.e., when g is true) or without error. Hence, X need not be concerned with the credal or expectation-determining probability that g is true in evaluating the rejection of g in relation to alternative expansion strategies.

Relative to the goal of obtaining new information without regard for truth value, the best potential answer is the rejection of all elements of U ; for that expansion strategy maximizes the information obtained as a result of expansion. But to do so is to expand deliberately into contradiction. The fact that no one is prepared to endorse this recommendation is some indication that the aim of an inquiry where the choice is between rival expansion strategies should not be to obtain new information regardless of truth value.

I have stated before that I take the position that the proximate aim should be to obtain new error-free information.

When X considers rejecting g , he is, as we have indicated, certain that he will obtain informational value worth $M(g)$. However, from his initial point of view (where he does not know whether g is true or false), there is a difference in the epistemic value or utility to X of rejecting g when g is true and when g is false.

According to proposals I have made elsewhere,⁸ the epistemic value of rejecting g can be represented by the number

$(1 - \alpha)M(g)$ when g is true and g is rejected erroneously and by $\alpha + (1 - \alpha)M(g)$ when g is false and rejecting g avoids error. The proposal is based on the idea that we can consider the effort to obtain error-free information as involving a balancing or trade-off between two desiderata: acquisition of new information and avoidance of error. If acquiring new information were the sole desideratum, options would be evaluated by reference to the values of the M -function. If avoidance of error were the sole desideratum, rejecting g would receive the value 1 if g is false and the value 0 if g is true. (Throughout this discussion, these "values" or "utilities" are to be understood to be unique up to a linear transformation so that in lieu of 1 and 0 any values a and b could have been assigned provided that a is greater than b . However, it is convenient to use 1 and 0. Similarly, instead of the function $M(g)$ any linear transformation of that function could be used. Once more convenience argues in favor of $M(g)$.)

In my view neither the desideratum of obtaining new information nor the desideratum of avoiding error should be favored to the exclusion of its rival. Hence, α should not be set at 0 or at 1. Indeed, α should never be less than .5; for if it is, rejecting g erroneously could bear higher epistemic utility than rejecting g' correctly where $M(g')$ is sufficiently lower than $M(g)$. The desirability of avoiding error should preclude this sort of evaluation.

2.6 Inferential Expansion

Criteria for evaluating expansion strategies to identify the best or a best strategy for the purpose of obtaining new error-free information should be derived by showing that these criteria pick out those strategies according to principles of rational choice which apply not only to such "cognitive decision problems" but apply also to problems where the aim is to choose an option (from a set of feasible options) in order to promote some practical, moral, economic, or political goal.

In this preliminary discussion, the general principle of rational choice I shall employ is the principle of maximizing expected utility. In chapter 4 and the chapters following, I shall replace this crude Bayesian decision theory with a more sophisticated approach which recognizes the applicability of the principle of maximizing expected utility in special cases but denies its applicability in general. The more sophisticated

decision theory entails modifications of the approach to evaluating expansion strategies to be outlined in this section. Nonetheless, the account offered in this section serves well as an introduction to the general approach to inferential expansion I favor. The modifications required according to the more sophisticated decision theory will not undermine the main ideas concerning inferential expansion outlined here. In any case, these ideas characterize the common thread connecting the address to inferential expansion advocated in this book and the proposals advanced in *Gambling with Truth*, "Information and Inference," and "Acceptance Revisited."⁹

Consider the expansion strategy involving the rejection of all and only the elements of U which are disjuncts in g . The epistemic utility of rejecting g without error is, according to the argument in section 2.5, $\alpha + (1 - \alpha)M(g)$. The epistemic utility of rejecting g erroneously is $(1 - \alpha)M(g)$.

Let the credal probability the agent X assigns to the hypothesis that g is false be representable by the real number $Q(\sim g) = 1 - Q(g)$ where $Q(g)$ is the credal probability X assigns the hypothesis g . The strategy of rejecting all elements of the ultimate partition which are alternatives in g bears an expected epistemic utility equal to $Q(\sim g)[\alpha + (1 - \alpha)M(g)] + Q(g)(1 - \alpha)M(g)$.

Let $q = (1 - \alpha)/\alpha$. Because α is restricted to values from .5 to 1, q ranges between 1 and 0. When $q = 1$, the relative importance attached to the desideratum that error be avoided is at a minimum. Anyone who adopts a value of α determining that value of q is prepared to exchange the maximum allowable risk of error for a given amount of informational value. As q declines, X will require a much-reduced risk of error before he will be prepared to expand to obtain a certain amount of informational value. As q increases, I shall say that X 's "degree of caution" decreases or, using an expression introduced by R. C. Jeffrey, his "degree of boldness" increases.

It can be shown that any expansion strategy which bears maximum expected utility rejects each element h_i of the ultimate partition such that $Q(h_i)$ is less than $qM(h_i)$ where $Q(h_i)$ is the credal or expectation-determining probability of h_i relative to X 's corpus $K_{X,t}$. Furthermore, any expansion strategy which rejects the h_i 's such that $Q(h_i)$ is less than $qM(h_i)$ and fails to reject those h_j 's where $Q(h_j)$ is greater than $qM(h_j)$ bear

maximum expected utility. I have argued elsewhere¹⁰ that X should pick the optimum expansion strategy that rejects the fewest elements of the ultimate partition. If this practice is followed, the ideas being outlined here lead to the following criterion for selecting an expansion strategy.

Given a corpus $K_{X,t}$, finite ultimate partition U , information-determining probability function M defined over the Boolean algebra of elements of U , an expectation-determining probability function Q defined over the same algebra, and an index of caution q , X should reject all and only those elements of U satisfying $Q(h_i) < qM(h_i)$.

2.7 Stable Acceptance and Caution

Sometimes X may regard two (or more) problems as meriting serious investigation where the best expansion strategy relative to one problem conflicts with the best expansion strategy relative to the other. In that case, X is obliged to take both problems into account in determining how to expand his corpus. Some consideration of how this may be done is provided in section 8.6. It is important to keep in mind, however, that the criterion for inferential expansion stated in the preceding section presupposes that there are no rival problems with conflicting demands under consideration.¹¹

Assuming such freedom from conflict, the result of expansion may be that X still suspends judgment between some elements of U and might consider further expansion based on his new corpus. The informational values for the unrejected elements of U will be modified. If g is the disjunction of unrejected elements of U and h_j is unrejected, the new M -value for h_j becomes $M(h_j)/M(g)$.

The Q -values for the surviving elements of the ultimate partition are also modified in a well-known way (which I will discuss subsequently). Given the old value of q , the rejection rule can be reapplied to the new case and sometimes will lead to the rejection of additional elements of the initial ultimate partition.

This process can be reiterated until it ceases to lead to the rejection of additional elements of the ultimate partition. When the ultimate partition is finite, there will always be a stage where reiteration yields no further information. Indeed, it can be shown that the iterative process will terminate at $q = 1$, leaving all and only those elements of U unrejected which

meet the condition that $Q(h_i)/M(h_i)$ is at a maximum.† When the M -values for all elements of the ultimate partition are equal, this means that those elements of the partition bearing maximum Q -values will survive rejection.

Some authors have found the procedure of reiterated application of the criterion for expansion objectionable. However, if X expands his corpus the first time, he no longer regards the rejected elements of the ultimate partition as serious possibilities. He has a new enlarged corpus which he is entitled to use as a standard for serious possibility in further efforts at expansion.

However, in many and, indeed, most contexts, the results of reiterated application of the criterion for expansion when $q = 1$ do appear to be presystematically objectionable.

For example, suppose that a coin known to be unbiased is tossed n times (where n is odd). X wishes to predict the relative frequency of heads in the n tosses. His ultimate partition consists of the $n + 1$ hypotheses as to the true relative frequency. If each of these hypotheses bears equal informational value, they will all receive equal M -values of $1/(n + 1)$. If $q = 1$, an element of the ultimate partition specifying that the relative frequency is r/n will be rejected on the first application if and only if $\binom{n}{r}(.5)^n$ is less than $1/(n + 1)$. By reiterating the procedure, X will be led to the conclusion that the relative frequency of heads will be exactly .5. Presystematically, this seems absurd. It may be plausible to conclude that the relative frequency is approximately 50% but that it is exactly 50% even in cases where n is very large seems much too strong.‡

† Let $\Sigma_1 Q$ be the sum of the Q -values of elements of the ultimate partition U surviving rejection on the first application of the rule when $q = 1$, and let $\Sigma_1 M$ be the sum of the M -values.

If h_i goes unrejected after the first expansion, its new Q -value is $Q_1(h_i) = Q(h_i)/\Sigma_1 Q$ and its new M -value is $M_1(h_i) = M(h_i)/\Sigma_1 M$. The condition under which survivors will avoid rejection on the second application of the rejection rule is that for every such h_i surviving rejection the first time, $Q_1(h_i) \geq M_1(h_i)$. This holds if and only if for every h_i unrejected the first time, $Q(h_i)/M(h_i) \geq \Sigma_1 Q/\Sigma_1 M = \text{constant}$.

More generally, if h_i survives rejection on the first k applications, it will survive the application $(k + 1)$ if and only if $Q_k(h_i) \geq M_k(h_i)$ if and only if $Q(h_i)/M(h_i) \geq \Sigma_k Q/\Sigma_k M$. If U is finite, the only elements surviving rejection at some stage are those for which $Q(h_i)/M(h_i)$ is a maximum.

‡ In my opinion, this point decisively undermines an approach favored by K. Lehrer. See Lehrer's "Induction, Consensus and Catastrophe," in *Local Induction* (edited by R. Bohdan, Reidel, 1975, pp. 131–132).

To avoid this awkward result, it seems plausible to require that, in general, q should be less than 1. However, there are circumstances where q might be allowed to go to 1; so that although normally q should be less than 1 (and, indeed, less than .5), it seems unwise to impose a blanket requirement that this always be so.

The criteria for expansion introduced in this and the preceding section have been formulated for cases where the ultimate partition is finite. Their extension to cases where it is infinite involve technical complications, some of which have been mentioned previously. However, it will be useful to indicate how these criteria can be extended to some special situations.

Suppose that X knows that a given coin has been tossed n times and has landed heads r times. X does not know what the chance or statistical probability p of obtaining heads on a single toss is, although he knows that the tosses are stochastically independent and that the chance p of heads is the same on each of the n tosses. He attempts to obtain information as to the true value of p on the basis of what he already knows, and he identifies as his ultimate partition all hypotheses specifying exact values for p . His ultimate partition is non-countably infinite.

I shall suppose, as would normally be plausible in such a case, that X regards each hypothesis specifying an exact value for p as informative as any other. This evaluation of informational value can be represented by the density function $m(p) = 1$ for all values of p between 0 and 1.

I shall also assume that X 's credal probability judgments are representable by a density function

$$f(p) = \frac{(n + 1)!}{r!(n - r)!} p^r (1 - p)^{n-r}.$$

In this case, the rules for expansion described previously can be extended so that an estimate of the exact value of p is rejected if and only if $f(p)$ is less than $qm(p)$. Moreover, the process can be reiterated. It can be shown that when $q = 1$, the reiterations must approach the result that the true value of p is r/n . This is the value for p for which $f(p)$ is a maximum. When q is less than 1, the conclusion warranted by reiteration is that the true value of p falls in an interval around r/n whose

width depends on the value of q and the value of n . The larger n , the narrower the interval. The smaller q , the wider the interval.

Suppose that q is fixed at some value less than 1 and presumably less than .5 and that n is large enough to warrant as a result of reiteration the conclusion that p is in the interval from $(r/n) - \epsilon$ to $(r/n) + \epsilon$ where ϵ is arbitrarily small. If ϵ is sufficiently small, X may be prepared to "throw caution to the winds" and let $q = 1$. Then he can reiterate further and reach an exact point estimate as to the value of p .

Thus, if X has found out in a large number of tosses that the coin has landed heads every time, he may be justified in concluding that the coin always lands heads on a toss. He can do so even though the hypothesis that $p = 1$ has an expectation-determining probability equal to 0.

This example shows how a hypothesis bearing 0 credal probability can, under appropriate circumstances, be added to a corpus of knowledge. One does not have to assign hypotheses positive credal probability in order for this to happen, as many authors including Hintikka, Jeffreys, and Shimony have supposed.¹²

The coin example illustrates one class of cases where the ultimate partition is infinite. There are others, but these latter involve further technical complications. For the present, they shall be ignored.†

2.8 Conclusion

The criteria for expansion outlined in this chapter require further elaboration and illustration. My immediate purpose has been, however, to cite enough detail to indicate the sense in which inferential expansion may be understood as involving a trade-off between risk of error and informational value.

The proposed criteria are heavily context dependent. They depend on X 's corpus and credal state (i.e., his judgments of credal probability). They also depend on his demands for information and the potential answers he has identified as means for gratifying these demands. They depend on his evaluations of the informational values of these potential answers

† Section 13.5 contains a brief explanation of how to handle cases where U is countably infinite or consists of hypotheses representable by points in an n -dimensional space.

and an index q representing the extent to which he is prepared to risk error for the sake of information.

Furthermore, the proposed criteria for inferential expansion are based on general criteria for rational choice to be applied not only in contexts where the aim is to obtain new error-free information but also where the objectives are more practical. Underlying the proposals made here are views concerning credal or expectation-determining probability, utility or value, and criteria for choosing between feasible options; these views will be developed in greater detail in later chapters.

Finally, the account of inferential expansion I have given is consistent with my advocacy of infallibilism. To be sure, prior to expansion, there is a risk, from X 's point of view, that the information to be added to his standard for serious possibility is false. Yet, sometimes X is justified in taking the risk. Once X has implemented the expansion strategy and taken the risk, he evaluates serious possibility according to a new standard relative to which the new information added is no longer possibly false.

3.1
Corrigibility

Critics of accumulation models of the growth of knowledge deny that expansion is the sole way in which knowledge is modified or improved. Theories are often discarded and replaced by others inconsistent with them. Emphasis on this correct observation lies at the heart of the critical philosophy of K. R. Popper and various political models of the growth of knowledge such as T. Kuhn's vision of normal change alternating with revolutionary change or P. K. Feyerabend's anarchistic conception of scientific inquiry.

How is it possible to concede the obvious and acknowledge the legitimacy of removing assumptions from a body of knowledge and, indeed, replacing them with conflicting assumptions, while insisting on the infallibility of knowledge and on the avoidance of error as a desideratum of scientific inquiry? How is infallibilism compatible with corrigibilism?

Those who sense the difficulty sometimes beg the question by presupposing that infallibility is equivalent to incorrigibility. Or they assume that necessity and possibility belong to hypotheses independent of X 's knowledge. Objections based on such assumptions need not disturb us to any greater degree in regard to contraction and replacement than they did in regard to expansion.

There are, however, more serious problems. From X 's point of view at t , all items in $K_{X,t}$ are true. X is certain at t that they are true. There is no serious possibility that they are false. For X to contract his corpus is for him to surrender error-free information. Replacement involves not only the abandonment of error-free information, but also the substitution of information that, from X 's point of view at t , is certainly and infallibly false. If X does take all items in his corpus to be infallibly true and seeks error-free information, it appears to be counterproductive for him to contract or to replace certainties with hypotheses he is certain are false. Counter to what I have claimed, infallibilism presupposes incorrigibilism.

3.2
Contraction

Consider first the problem of contraction. X cannot import error into his corpus by contracting. Importation of error involves adding information to the corpus. Hence, the desideratum of avoiding error presents no obstacle to contraction.

The trouble with contraction is that it entails a deliberate loss of information which, from X 's initial point of view, counts as certainly and infallibly true. Since X is presumably interested in obtaining information, what could induce him to contract and lose information? If X thought there was a risk of error in retaining the information whose removal is under consideration, there might be an inducement to contract. But, under infallibilism, X should not, from his point of view, regard himself as running such a risk. The loss of information through contraction appears gratuitously counterproductive.

There are cases, however, where contraction is responsive to a legitimate demand. If X detects inconsistency in his initial corpus $K_{X,t}$, he has excellent reason to contract. An inconsistent corpus fails as a standard for serious possibility to be used in inquiry and deliberation. The corpus is useless to X at t and should be modified.†

† When X 's corpus expressible in L is inconsistent, that corpus breaks down as a standard for assessing the possible truth of sentences in L and for specifying truth conditions for sentences in L . This may seem to provide solid grounds for prohibiting X 's corpus from ever being inconsistent—at least if he is rational. However, as we have noted, even when X has perfect memory and suffers no limitations in his capacities to make calculations and draw explicitly the deductive implications of his assumptions, his corpus can become inconsistent due to routine expansion. We should not talk of prohibiting inconsistent corpora but of the problem of shifting from an inconsistent corpus via contraction. The breakdown of the functions which corpora of knowledge are intended to perform in inquiry and deliberation when they are inconsistent furnishes good reason for undertaking efforts to shift from them; but this should not be confused with good reasons for prohibiting inconsistency in the first place.

When X shifts from corpus K_1 to corpus K_2 , any justification X can offer to himself for making the shift should be based on the assumption (expressed in the metalanguage L_1) that all items in K_1 are true (in L) and infallibly so. However, when K_1 is inconsistent, X cannot proceed in this manner. If he did so, the inconsistency infecting K_1 would spread to his metacorporexpressible in L_1 , and so on. X would have no coherent basis for evaluating alternative ways to contract from K_1 .

For this reason, I suggest that we look on such situations as cases where X treats the object language L syntactically (as long as his corpus in L is inconsistent) and treats K_1 and other potential contractions of K_1 as so many different uninterpreted systems of sentences. He can, in this way, retain the consistency of his metacorporexpressible and avoid begging questions as to which of the

Even if *X* were so ideally situated and competent that he avoided mistakes of computation and memory and, hence, did not contradict himself due to lapses of these kinds, he could still end up with an inconsistent corpus. *X* cannot justifiably expand his corpus into contradiction via inferential expansion. Doing so would deliberately import error into *X*'s corpus. But contradiction can be injected via routine expansion, as explained in the previous chapter. Observations sometimes contradict our most cherished theories. Men whose authority we respect sometimes disagree with us. In both cases, the result is pressure to contract.

Other good reasons exist for contracting a corpus. Suppose the initial corpus contains some theory T_1 . A second theory T_2 contradicts T_1 . From *X*'s initial point of view, T_2 is certainly false. Yet it may be superior in all other respects to T_1 as a means for furnishing systematic explanations in some domain. *X* could recognize the superior explanatory virtues of statistical mechanics even though he is certain that it is false and that classical thermodynamics is true.

In such cases, *X* might be prepared to suffer a loss of information due to the removal of T_1 from his corpus in order to be in a position to take the truth of T_2 to be seriously possible. In that event, *X* exchanges a given amount of information for the opportunity to give an informationally attractive hypothesis a hearing.

To contract in this case is not to reject T_1 as false and to accept T_2 as true but to shift to a position where judgment is suspended between these rival hypotheses so that investigations can be undertaken to decide whether T_1 should be reinstated via inferential expansion or T_2 should take T_1 's place.

Notice that contraction to give some new theory a hearing is by no means automatic. A case must be made for giving the new theory a hearing by establishing its informational value.

potential contraction strategies to adopt by assuming at the outset that some items in K_1 are possibly false and others are not. Of course, sentences in UK will, on my view, be immune from revision, and discriminations will be made between other items in K_1 with respect to vulnerability to revision. But these discriminations will not involve any presupposition that the items less vulnerable to revision are more likely to be true than those which are more vulnerable.

3.3 Degrees of Corrigibility

Given a demand for contraction occasioned either by the need to remove contradiction or to give a hypothesis a hearing, how should a contraction strategy be chosen? In general, several such strategies will be available; any one of them can provide the desired hearing.

Some potential contraction strategies will be superior to others—at least relative to the problem which occasioned the demand for contraction. Hence, relative to that problem, some items in the corpus prior to contraction will be more vulnerable to removal than others.

Thus, sentences in a corpus will differ from one another with respect to grades of corrigibility or vulnerability to removal from the corpus.

It would be a mistake, however, to construe such grades of corrigibility as grades of certainty or probability. This widespread error is nothing but a manifestation of the confusion of infallibility and certainty with incorrigibility.

From *X*'s point of view prior to contraction, all assumptions in his corpus are equally certain and infallible.† Yet some are more vulnerable to removal than others. Some assumptions are maximally certain and infallible, and, yet, are highly corrigible. Others may be equally as certain and yet eminently incorrigible (e.g., items in the urcorpus). Grades of corrigibility cannot be grades of certainty.‡

In contraction, the aim should be to minimize the loss of informational value resulting from contraction subject to the constraint that the need occasioning the demand for contrac-

† When the initial corpus to be contracted is consistent, all items in the corpus are maximally certain and are infallible—from *X*'s point of view. When the initial corpus is inconsistent, this remains the case from *X*'s point of view: but, from *X*'s point of view, the negation also holds. The point is that an inconsistent corpus no longer can serve as a standard for serious possibility. On the other hand, as pointed out in the preceding footnote, no relevant difference can be recognized between items in the inconsistent corpus with respect to certainty and infallibility. Yet, some items are more vulnerable to removal than others.

‡ We may go further. Distinctions between statements with respect to corrigibility can be drawn only between statements which are members of a corpus and, hence, are all equally and maximally certain relative to that corpus. If both the truth and falsity of h is a serious possibility for *X* at t and, hence, *X* assigns h an intermediate grade of credal probability between 0 and 1, it will make no sense for *X* to also evaluate h with respect to its vulnerability to removal. From what corpus is it to be removed?

tion in the first place be satisfied. If this view is sound, corrigibility decreases with informational value. X should be prepared to remove a relatively uninformative assumption from his corpus rather than remove an assumption which will seriously weaken that corpus as a standard for serious possibility.

In cases where the losses of informational value resulting from two or more contraction strategies that satisfy the need occasioning the demand for information are all minimal relative to the set of all such contraction strategies identified by the X , then X has no basis for deciding between them. Rather than choose one arbitrarily, he should, if feasible, suspend judgment by implementing all of these optimal contraction strategies jointly.

Thus, suppose a theory firmly ensconced in the corpus K_1 is contradicted by new data; so that X shifts to a new corpus K_2 that is inconsistent due to the presence of the theory and the data. X could remove the theory and retain the observation reports. Or he might remove the observation reports and retain the theory. In doing the latter, he might be led to question the reliability of the routine for observation which, as a result of contraction, will be judged to have bred false reports. Thus, sometimes the loss of information suffered by endorsing one contraction strategy may be roughly the same as that resulting from adopting the other. In that case, X should remove both the theory and the observation reports from his initial corpus K_1 and check on both the theory and the observation reports.

3.4 Contradiction vs. Anomaly

The conflict of data with an already-established theory should not be confused with the conflict of data with a theory which is a potential answer not yet incorporated into the body of knowledge. In the latter case, the verdict of the data tends to be decisive and the conclusion that the theory is false is added to the corpus.

When the theory is not a hypothesis but a settled assumption, the conflict cannot be construed as the outcome of the test of the theory. Assumptions taken for granted as knowledge are not tested while they are so taken. Only when they are removed from the corpus and cast into doubt will they be tested.

Michelson's first experiments were not conducted to test Newtonian mechanics, Maxwell's electromagnetic theory, or

the ether hypothesis. Michelson took these for granted as background knowledge, and was attempting to test only a hypothesis of Stokes concerning the relative motion of the earth and the surrounding ether.¹ When Lorenz subsequently noticed that the Stokes hypothesis conflicted with mechanics and electromagnetic theory, Michelson's results were recognized to have generated a conflict within the settled scientific corpus. Not only were various items in that corpus subjected to scrutiny, but so were the data obtained by Michelson.

Sometimes the results of observation present anomalies for settled theories. Anomalous data do not contradict the settled theories. They are phenomena which the settled theories allegedly should be able to explain (perhaps supplemented with suitable collateral assumptions), but for which no explanations, in terms of these theories, have thus far been found.

The difference between anomaly and contradiction is important. When data contradict a settled theory, the demand for contraction is always legitimate. When data present anomalies for settled theories, the legitimacy of demands for contraction is by no means automatic. Indeed, the thrust of research is to find explanations of the anomalies which eliminate their anomalous character with the aid of the settled theories. When this is done successfully, contraction may be avoided altogether. Even when such efforts have failed, contraction will not become legitimate unless a rival theory can be identified relative to which the anomaly is explicable. Thus, anomaly generates pressure to contract only when a theory inconsistent with what is currently in X 's corpus furnishes explanations more satisfactory than those available through the resources of the corpus itself. Contraction is then warranted in order to give an informationally attractive theory a serious hearing.

3.5 Replacement

The most impressive revisions of scientific knowledge (many of which are called "revolutions") are replacements. Replacements are shifts from corpora to other corpora inconsistent with the initial ones.

From X 's initial point of view when he adopts a theory T_1 , replacement of T_1 with T_2 inconsistent with T_1 is tantamount to the deliberate substitution of a certainly false hypothesis for one which is certainly true. That is to say, that is the way

X should view the matter if his corpus serves as his standard for serious possibility and the theses of epistemological and categorical infallibility obtain. If, in addition, *X* intends to avoid error in revising his corpus, it is utterly counterproductive from his point of view to undertake the replacement.

The legitimacy of replacement could be denied; but this would be to condemn a historically important mode of development of scientific knowledge as illegitimate.

One could deny the relevance of avoidance of error in justifying replacements; but to do so is incompatible with my thesis that avoidance of error is an invariant feature of the proximate aims of efforts to revise a body of knowledge.

Categorical and epistemological fallibilism could be endorsed; but this is also incompatible with the epistemological outlook of this book.

There is, however, a fourth alternative. We might deny the legitimacy of efforts to justify shifting directly from K_1 containing *h* to K_2 containing $\sim h$, but might allow the legitimacy of such replacement provided that it can be decomposed into a sequence of contractions and expansions each of which is justified.

Suppose that at t_1 when he adopts K_1 , *X* has good reason to contract his corpus by removing *h* in order to give *g* inconsistent with *h* a hearing. The result of contraction will be a shift to K_3 relative to which *X* will suspend judgment between *h* and *g* and, hence, between *h* and $\sim h$.

X may then engage in inquiry in order to determine whether to expand by readmitting *h* into his corpus or by adding *g* instead. Often such investigation will involve obtaining new information via observation or other routine; so that *X* might shift to corpus K_4 containing more data, but which as yet does not settle the question regarding *h* and *g*. But whether *X* does shift from K_3 to K_4 through subsidiary investigation, he will often stop and, relative to what he then knows, will contemplate expansion in order to answer the question.

If *X* expands by adding *g* and, hence, shifting to corpus K_2 , the net effect is a shift from K_1 to K_2 inconsistent with K_1 —i.e., replacement. Yet, the replacement may be decomposed into a sequence of revisions each of which is either a contraction or an expansion justified in a manner which respects avoidance of error as a desideratum in inquiry.

The key idea here is scarcely novel. When *X* considers replacing *h* with $\sim h$, if he attempts to settle the matter from his initial point of view, he will beg the issue in favor of *h*. If he considers by way of anticipation how matters would appear once he has endorsed $\sim h$, the question will be begged the other way. By contracting to a corpus containing neither *h* nor $\sim h$, he may explore the matter without committing himself one way or the other to the question under dispute.

Notice that residual shifts may also be analysed as sequences of contractions and expansions. Hence both contraction and expansion can be viewed as the fundamental types of revision subject to critical control, and all other sorts of revisions may be then understood as sequences of changes of these kinds.

I am not claiming that the historical record will reveal that replacements of one theory by another always take place as the net result of an explicitly or consciously implemented sequence of contractions and expansions. However, if such replacements are defensible, they should be decomposable (for purposes of analysis) into sequences of this sort; in such a sequence, each step must be justifiable.

3.6 Contextualism

As I understand him, Kuhn claims that there are at least some occasions in the development of science when replacements cannot be rationalized along the lines I am suggesting, but where their scientific legitimacy should nonetheless be acknowledged. However, there is one important respect in which my approach is closer to that of Kuhn than it is to that of many of his critics.

The benefits we have reaped from the legacies of Frege and Husserl have been accompanied by the curse of an exaggerated hostility to psychologism. The lust for "objectivity" has led many sober authors to equate the context-dependent with the idiosyncratic and both with the "irrational," "arbitrary," or "subjective." Consequently, insofar as what counts as legitimate revision of the state of human knowledge is taken to be context dependent (i.e., dependent on the research programs of the investigators, the demands for information engendered by these programs, the problems encountered, the potential answers identified as solutions to these problems, etc.), it is taken to be beyond systematic critical control and

to be "subjective" or "irrational." Only historians, psychologists, sociologists, and other social scientists can reflect upon it with profit.

Kuhn has explicitly denied that the "growth of knowledge" is beyond critical control; yet he has emphatically insisted on the "historical"—i.e., context-dependent—character of scientific change. By implication, he must be rejecting the assumption that what is context dependent is beyond critical control. Indeed, his rejection is rather explicit.

What I am denying then is neither the existence of good reasons nor that these reasons are of the sort usually described. I am, however, insisting that such reasons constitute values to be used in making choices rather than rules of choice. Scientists who share them may nevertheless make different choices in the same concrete situation. Two factors are deeply involved. First, in many concrete situations, different values, though all constitutive of good reasons, dictate different conclusions, different choices. In such cases of value-conflict (e.g., one theory is simpler but the other is more accurate) the relative weight placed on different values by different individuals can play a decisive role in individual choice. More important, though scientists share these values and must continue to do so if science is to survive, they do not all apply them in the same way. Simplicity, scope, fruitfulness, and even accuracy can be judged quite differently . . . by different people.²

Everyone (or nearly everyone) agrees that if X and Y have different bodies of evidence, they may, as reasonable men, disagree in the choices they make or the conclusions they reach. To this extent, even anticontextualists are willing to allow for the relevance of context in formulating critical standards.

Kuhn points out that even if X and Y agree on the evidence but differ in what they regard as valuable in a scientific hypothesis, they might legitimately come to different conclusions. Indeed, even if they agree on what are the marks of an attractive hypothesis (e.g., if they regard simplicity and accuracy as such desiderata), they might differ in their judgments as to how these desiderata are to be weighed.

On these matters, I have no quarrel with Kuhn. Thus, in my account of inferential expansion, room is given for agents to differ not only with respect to an initial body of knowledge,

but also with respect (i) to the potential answers identified, (ii) to evaluations of such answers with respect to informational value (as represented by an M-function) and (iii) to the relative importance attached to avoidance of error as compared to the quest for new information as mirrored in the index of caution (section 2.5). Other contextual factors are relevant to the appraisal of contractions. Nonetheless, taking context seriously in appraising expansions and contractions does not imply that revisions of these kinds are immune from critical evaluation. It means only that relevant contextual parameters have to be taken into account in such evaluation.

Suffering from an overdose of antipsychologism, followers of R. Carnap and Popper often overlook this point. They charge that those who emphasize the context-dependent character of human knowledge are focusing on the irrational and idiosyncratic at the expense of the rational and objective. Kuhn maintains there are other alternatives. So do I.

3.7 Revolutionary vs. Normal Science

But Kuhn defends another dualism which I find untenable—namely, the contrast between revolutionary and normal science. Although Kuhn does not explicitly say so, replacements which occur in the course of normal science may be reconstructed as sequences of contractions and expansions along the lines I have outlined. One can move to a neutral basis relative to which no questions are begged and adjudicate the merits of rival hypotheses in a manner which respects the desirability of avoiding error and taking risk of error into account.

In a revolutionary phase, replacement of one theory T_1 by another T_2 cannot proceed in this manner. Contracting from the initial corpus T_1 to a corpus which begs no questions concerning the truth values of T_1 and T_2 cannot happen. The contracted corpus cannot serve as a standard for serious possibility and as a basis for a truth definition. Consequently, the contracted corpus is useless for the purpose of evaluating the rival hypotheses in a neutral manner and, in particular, for taking avoidance of error into account in subsequent efforts at expansion. Whatever else Kuhn might mean by calling two such theories T_1 and T_2 "incommensurable," he does appear to be quite clear that incommensurability precludes moving to a neutral non-question-begging base and, hence, in my jargon,

precludes rationalizing a replacement as a contraction followed by an expansion.

Thus, Kuhn writes: "Sir Karl takes it for granted that proponents of competing theories do share a neutral language adequate to the comparison of such observation reports. I am about to argue that they do not. If I am right, then 'truth' may, like 'proof,' be a term with only intra-theoretic applications."³

Thus, revolutionary changes in Kuhn's sense appear to be akin to what I take to be changes in conceptual framework due to revisions in what counts as conceptually but not categorically incorrigible. I have already stated my aim to construct an account of the revision of knowledge which denies that there should be revisions of this sort. On my view, there are no revolutionary changes or, at any rate, there should not be.

On the other hand, this does not mean that I advocate normal science in a sense akin to Kuhn's.

In the first place, although I am inclined to the view that there are some fixed methodological norms (counter to the position of Feyerabend⁴), these norms are extremely weak and heavily dependent for their operation on contextual factors which change with developments in scientific inquiry. Consequently, I too am "against method" if by this one means a very substantive method immune to criticism and revision during the course of inquiry itself. Nonetheless, I remain convinced that at any given stage of inquiry, Feyerabend's insistence that "anything goes" should be rejected.

Secondly, even though I contend that we should judge truth earnestly relative to our evolving doctrine, I do not mean to suggest that institutional arrangements for the conduct of inquiry should be set up so as to suppress the development of rivals to currently established doctrine and method. In section 1.10, I pointed out that my rejection of fallibilism need not undermine promotion of institutions protecting free speech. Similarly nothing in my insistence that all legitimate revisions of corpus should be justified expansions or contractions which respect the desirability of avoiding error and the acquisition of worthwhile information should be construed as entailing curtailment of imaginative efforts to construct rivals to re-

ceived doctrine attractive enough to warrant contraction in order to give them a hearing.

To be sure, I remain unconvinced by Feyerabend's contention that positive encouragement of the proliferation of rivals to settled doctrine is desirable for progress in inquiry or the development of authentic, independent agents. I grant that the construction of rival theories is crucial to the successful appeal to anomaly as a justification for removing a settled theory from its privileged status as a part of a standard for serious possibility. When earnest efforts to accommodate anomaly utilizing existing theory have failed, it makes sense to encourage the construction of rival theories which do remove the anomaly; for then the inquirer may be in a position to justify contraction of his corpus to give a hearing to an alternative which promises to improve the informational value of his corpus. On this basis, a good case can be made for protecting the rights of inquirers to proliferate rivals. But it does not follow from this concession that the proliferation of rival theories should be positively encouraged regardless of whether the need to accommodate anomaly by appeal to alternatives to existing doctrine is urgent or not.

Feyerabend would, of course, dismiss this advocacy of judicious proliferation. On his view, we do not proliferate rivals to provide good reasons for contraction to give hypotheses which promise to accommodate anomaly better than current doctrine a proper hearing. Feyerabend, like Kuhn, thinks that rivals to comprehensive theories are incommensurable so that the idea of giving a rival a hearing makes little sense. A shift from a corpus containing one comprehensive theory to a corpus containing a rival theory is always a replacement where the desideratum of avoiding error is flouted. Yet, Feyerabend believes that it is a good thing for such replacements to take place and that proliferating rivals is an effective means for doing so.

I mean to reject Feyerabend's cognitive anarchism or da-datism just as I reject Kuhn's vision of revolution alternating with normal science; and I reject the thesis of incommensurability which is central to both of their outlooks.

Kuhn and Feyerabend supply many illustrations of incommensurability. I believe these examples can be reconstructed so as to disarm the claims: but I shall not try to substantiate

my view here. Instead, I shall attempt to provide the fragments of a positive account of the revision of knowledge predicated on the assumption that all legitimate revisions of knowledge may be rationalized as sequences of contractions and expansions. There are no incommensurable theories. There are no revolutionary replacements.

3.8 Myopia and Convergence to the Truth

Popper, as C. S. Peirce before him, emphasizes that approaching the true complete story of the world is the ultimate long-run aim of scientific inquiry. Kuhn raises doubts about this view.⁵ So do I. Those who sympathize with some variant of the Peirce-Popper view must find my position as objectionable as they find Kuhn's. Indeed, they must find replacements of allegedly incommensurable theories immune to rationalization along the lines I have suggested illegitimate; but they must also object to replacements where the rival theories are commensurable and the replacements can be rationalized as sequences of contractions and expansions.

Consider the charge that my account of replacement is pedantic and formal sham. When X shifts from K_1 to K_3 by removing h and giving g a hearing, he will often be able to calculate that from his new position he will be justified in expanding by adding g and, hence, $\sim h$ to his corpus K_3 . Thus, prior to contraction X may be certain that if he contracts he will subsequently add $\sim h$ to his corpus so that the net effect will be to substitute something he is certain is false (namely, $\sim h$) for something he is certain is true. In such cases, contraction courts error just as effectively as does replacement!

It is no use denying that situations of this sort arise. If X contemplates removing classical thermodynamics from his corpus after finding out the results of the experiments of Svedberg and Perrin on Brownian motion, he may anticipate that he will then be justified in endorsing statistical mechanics. If prior to contraction he is certain that statistical mechanics is false, he would be foolish to contract if he were seriously concerned to avoid error in subsequent revisions of his corpus.

In any case, even if prior to contraction X is not certain that removing h will lead to his accepting g into his corpus, there will, from X 's point of view, be some positive probability that he will do so. Hence, counter to what I suggested previously, even contraction seems to incur a risk of error—error that,

when combined with the loss of information it also incurs, renders it extremely difficult to compute the costs and benefits of contraction and subsequent replacement.

This objection rests on the following assumptions: (a) the agent X uses his initial corpus as his standard for serious possibility and as the basis for the conception of truth he employs in seeking to avoid error and (b) in seeking to avoid error as understood relative to his initial corpus, the agent should be concerned to avoid error not only at the next step but in subsequent revisions as well.

Peirce and Popper agree that the ultimate aim of inquiry is convergence on the truth—i.e., the true complete story of the world. As I understand their views, this means that if they thought that (a) obtained, they would also embrace (b). But then, for reasons just explained, contraction would become indefensible. Since Peirce and Popper both endorse the revisability of knowledge, they are committed by their conception of the aims of inquiry to reject (a). But this means they reject the thesis of epistemological infallibilism, and thus the concept of knowledge as a standard for serious possibility which entails that thesis.

I agree with Feyerabend, Kuhn, and W. V. O. Quine in rejecting the Peirce-Popper concept of the ultimate aims of scientific inquiry. I do not deny that X may have longer-run objectives than those impinging directly on the revision immediately under consideration. But I do deny that truth in the sense explained plays a significant role in such aims.

On the other hand, I do claim, counter both to Feyerabend-Kuhn-Quine and Peirce-Popper, that avoidance of error is an invariant feature of the diverse proximate goals directing specific inquiries concerned with the revision of knowledge. In contemplating the contraction from K_1 to K_3 , X should be concerned as to whether in taking that step he will import error into his corpus. He should not be concerned with the prospects of doing so at subsequent stages. However, once he does contract and thus shifts to K_3 , he should be concerned to avoid error as judged relative to K_3 in contemplating further expansions or contractions of K_3 . In this respect, X 's concern for truth and avoidance of error should be myopic.

I cannot prove that I am right in singing the praises of myopia. Nonetheless, the alternatives seem far less attractive.

According to the Peirce-Popper view of the ultimate aims of inquiry, the revisability of knowledge via contraction and replacement can be legitimate only if knowledge does not serve as a standard for serious possibility and infallibilism is abandoned.

But if one embraces fallibilism, it becomes a mystery as to why one should seek to revise knowledge; for one is then entitled to ask whether knowledge has any use at all, either practical or theoretical.

A possible way out is to insist on a double standard for serious possibility: one for practical purposes and one for theoretical or cognitive purposes. I doubt whether this view is coherent. In any case, it implies an untenable dualism between theory and practice.

Finally, although both Peirce and Popper insist on convergence to the truth as characteristic of the aims of scientific inquiry, neither author explains how avoidance of error can function as a desideratum in the resolution of the special problems of specific inquiries. Their concern with truth is an empty act of piety. For them, the avoidance of error cannot play a role as a desideratum of the proximate aims of specific scientific inquiries (for both normal or revolutionary inquiries); and truth turns out to be irrelevant in precisely the context where it ought to count the most.

This is the kind of view one is led to if one eschews myopia and seeks convergence to the truth as an ultimate aim of inquiry. In my opinion, myopia wins hands down!

3.9 Theoretical and Practical

I have been trying to outline a concept of the uses of knowledge in inquiry and deliberation, a view of the proximate aims of efforts to improve knowledge, and an explanation of how, given such aims, one can consistently claim that knowledge is corrigible and yet infallible.

I have treated the various modes of revising knowledge on the assumption that revising knowledge is a goal-directed activity and that revisions of knowledge should be evaluated in terms of estimates of their efficacy in realizing given cognitive goals.

I also believe that the principles of rational choice or rational goal attainment governing deliberation in science ought to be the same as those regulating the rational attainment of moral,

political, economic, and other practical objectives. The differences between theoretical inquiry and practical deliberation is a difference in goals and not a difference in the criteria for rational choice that regulate efforts to realize these goals.

Thus, even an outline of an account of the improvement of knowledge of the sort advanced here should offer some system of principles for rational choice. It should also furnish some characterization of the way probability judgment cooperates with knowledge and value judgment in guiding rational choice. Probability is important in other ways as well. Thus concern for avoiding error manifests itself in inquiry as a concern to minimize risk of error. But risk of error is best understood in terms of probability.

Because probability judgment supplements *X*'s corpus of knowledge as a resource for inquiry and deliberation, a systematic account of the improvement of knowledge which neglects the problem of improving probability judgment is seriously defective.

For these reasons, the remainder of this book will be given over to a discussion of rational choice, the role of probability judgment in decision making, and an account of the improvement of probability judgment.

I do not claim that this volume contains a complete account of these matters. To the contrary, I aim merely to outline a framework within which more detailed investigations may be pursued. I hope, however, that enough will have been said to delineate the epistemological program I am advocating.

4.1 The Four Questions

Whoever denies he is certain of anything is neurotic. Whoever claims certainty of everything is foolish. For sufficiently rich L , X 's corpus expressible in L should rarely, if ever, be a maximally consistent set. X should be certain of the truth values of some hypotheses and uncertain about others.

The corpus $K_{X,t}$ allows X to discriminate at t between what is and what is not seriously possible but by itself fails to grade serious possibilities with respect to probability.

Yet, such discriminations are relevant to deliberation and inquiry. Their relevance is best explained in terms of an account of rational choice that shows how judgments of credal probability, together with the underlying evaluations of hypotheses with respect to serious possibility, contribute to the evaluation of rival feasible options in decision making.

Such an account of rational decision making is important for another reason. X 's corpus of knowledge (i.e., his standard for serious possibility) and his credal state (i.e., his system of judgments of hypotheses with respect to credal probability) constitute X 's cognitive resources for deliberation and inquiry. A comprehensive approach to the improvement of knowledge should seek to account for the improvement of all cognitive resources. According to the outlook expressed in the previous chapters, such improvements are judged such because they are taken to be admissible means for promoting the objectives which prompt efforts to improve cognitive resources. The criteria for admissibility, however, should be sufficiently general to be applicable to all forms of deliberate decision making, regardless of whether the options and the aims are "cognitive" or whether they are moral, political, economic, or otherwise "practical."

Thus, an account of the principles of rational decision making is fundamental to an approach to the improvement of cognitive resources of the sort I seek to develop; not only

because it illuminates the sense in which these cognitive resources are, indeed, resources for deliberation and inquiry, but also because it should constitute a cornerstone for constructing principles for evaluating any subsequent revisions of these resources.

An account of deliberate rational decision making furnishes criteria for appraising the admissibility of feasible options relative to X 's cognitive resources—i.e., his corpus of knowledge and credal state—and relative to X 's system of valuations induced by his goals and values. Consequently, it will be necessary to discuss general conditions of rational valuation, preference, or utility. Conditions on potential corpora and, hence, on the evaluation of hypotheses with respect to serious possibility have already been considered. But attention will have to be given to conditions for the rational evaluation of hypothesis with respect to credal probability.

In this book I shall not explore criteria for evaluating would-be improvements in all factors which are resources for deliberation and inquiry, but only improvements in those factors which I have called cognitive resources—i.e., corpora of knowledge and credal states. Improvements in goals and values will not be considered in any systematic manner.

Thus, we have four questions to consider:

- (a) *The question of rational cognition:*
 - (i) *Rational standards for serious possibility:* The specification of conditions on the evaluation of hypotheses with respect to serious possibility leading to the characterization of potential corpora given in chapter 1.
 - (ii) *Rational credence:* The specification of conditions on credal states which ought to be satisfied by all rational agents insofar as they are able.
- (b) *The question of revising cognitive resources:*
 - (i) *Revising corpora of knowledge.*
 - (ii) *Revising credal states.*
- (c) *The question of rational valuation.*
- (d) *The question of rational choice.*

Considerable attention has already been devoted in the first three chapters to part (i) of both (a) and (b), and these topics will reappear from time to time in the remainder of this book.

Far less space has been devoted to the other two parts of the first two questions. This lack will be more than compensated for in the chapters that follow. And questions (c) and (d) will be examined to an extent sufficient to provide some foundation for the answers offered to both parts of the first two questions.

The present chapter outlines my general approach to answering these questions. The rest of the book contains qualifications, elaborations, and justifications of various components of that outlook. Some of the subsequent commentary amounts to technical fine tuning; but much of the later discussion, so I think, is critical to an understanding of the epistemological stance I am adopting. I do not seek in this chapter to summarize my viewpoint, but to offer a guide to the reader seeking to find his way through the complex argument that follows.

4.2 Rational Credece

Let K be a consistent potential corpus expressible in L . A real-valued function $Q(h; e)$ defined for all h in L and all e in a subset E of L is a *finitely additive and normalized probability measure in L relative to K* if and only if $Q(h; e)$ satisfies the following conditions:

- (1) Let E be the set of all hypotheses in L consistent with K . If $e \in E$, then $Q(h; e) \geq 0$.
- (2) If $K \vdash h \equiv h'$ and $K \vdash e \equiv e'$, then $Q(h; e) = Q(h'; e')$.
- (3) If $e \vdash \sim(h \& g)$, then $Q(h; e) + Q(g; e) = Q(h \vee g; e)$ (*additivity*).
- (4) If $K, e \vdash h$, then $Q(h; e) = 1$ (*normalization*).
- (5) $Q(h \& f; e) = Q(h; f \& e)Q(f; e)$ (*multiplication theorem*).

When $K \vdash e$, then $Q(h; e) = Q(h)$ states the *unconditional probability* that h or the unconditional Q -value for h .

A credal state B relative to a potential corpus K is a way of evaluating hypotheses with respect to credal probability relative to the system of valuations with respect to serious possibility induced by adopting K as the standard for serious possibility. I shall represent such a credal state as a set B of real-valued functions $Q(h; e)$. Conditions of credal rationality will be formulated as conditions on such sets and their elements. The motivation and import of this method of repre-

senting credal states will emerge as the discussion develops in this chapter and in the remaining portions of this book.

Suppose $K \vdash h \equiv h'$. If X should adopt K as his standard for serious possibility, he should evaluate h and h' in exactly the same way with respect to serious possibility. Either both are evaluated as possible or as not possible. Similarly, both hypotheses should be evaluated in the same way with respect to credal probability. I take it that anyone who is willing to countenance judgments of subjective or credal probability will concede that ideally situated rational agents should make judgments of credal probability conforming to this requirement.

This condition, however, suffices to show that if X adopts K as his standard for serious possibility, certain credal states are forbidden to him which might be permitted had he endorsed some alternative corpus.

Thus, conditions for credal rationality must be formulated relative to potential corpora and may be thought of as conditions on a pair (K, B) where K is a potential corpus expressible in L and is deductively closed and contains UK .

Credal Coherence: If $Q \in B$, then Q is a finitely additive and normalized probability measure relative to K .

Credal coherence is an example of what I shall call a principle of *inductive logic*. Whether agent X endorses corpus K or not and regardless of X 's other circumstances or who he or it is, X should regard any credal state relative to K which contains a Q -function violating credal coherence as forbidden to rational agents who endorse K . Any such Q -function is *logically impermissible relative to K* .

Some authors who endorse credal coherence would favor strengthening the requirement by insisting that E be restricted to the set of hypotheses bearing positive Q -values in condition (1) on probability measures, that all hypotheses consistent with K and e bear positive Q -value conditional on e (*credal regularity*) and that condition (3) be supplemented by a principle of *countable additivity*.

I shall not impose any of these conditions, and this decision will prove relevant to some of the technical developments in the subsequent discussion. These matters will be discussed in chapter 5, and sections 12.14-12.15 and 15.9.

B. De Finetti¹ and L. J. Savage² contend that credal coherence is the sole principle of inductive logic. Other authors have endorsed much more powerful inductive logics. I shall return to this controversy later in this chapter and elsewhere in this book.

Some principles of credal rationality do not qualify as principles of inductive logic because they do not unconditionally prohibit membership to Q in B relative to K . Rather they specify how many Q -functions satisfying the requirements of inductive logic may be in a single set B . I endorse the following two conditions:

Credal Consistency: $B \neq \emptyset$ if and only if K is consistent.

Credal Convexity: For any Q -function relative to K and e consistent with K , let $Q_e(h) = Q(h; e)$. Let B be a credal state relative to K . B_e is the set of all Q_e -functions obtained from Q -functions in B . For every e consistent with K , B_e is convex—i.e., if $Q_e^1 \in B_e$ and $Q_e^2 \in B_e$, then $Q_e^\alpha \in B_e$ where $Q_e^\alpha = \alpha Q_e^1 + (1 - \alpha)Q_e^2$ for every α such that $0 \leq \alpha \leq 1$.

Many philosophers, statisticians, and decision theorists advocate a view which I shall call *strict Bayesianism*. Advocates of strict Bayesianism offer characteristic answers to the four questions raised in section 4.1. To be sure, H. Jeffreys³ and R. Carnap⁴ differ not only with each other but with the views of De Finetti,⁵ and Savage⁶ (and the later Carnap).⁷ Nonetheless, disputes among strict Bayesians arise within a framework of broadly shared agreements which furnish comprehensive guidelines concerning how the four questions ought to be answered.

In particular, it is worth noting that strict Bayesians all agree in endorsing the conditions of credal coherence, credal consistency, and credal convexity. To this extent, the view of credal rationality I have been developing should prove quite acceptable to the most orthodox Bayesians.

To be sure, there will be differences concerning the propriety of strengthening credal coherence or regarding the merits of adding further principles of inductive logic. But disagreements concerning these matters already exist among strict Bayesians.

The point where my view departs from strict Bayesian or-

thodoxy concerns credal convexity. Strict Bayesians endorse a stronger condition than credal convexity:

Credal Uniqueness: There is at most one Q -function in B .

I reject credal uniqueness as a condition on credal states obligatory on all rational agents under all circumstances. Rational X 's credal state need not be representable by a single Q -function.

Rejection of this fundamental tenet of strict Bayesian doctrine is not new. J. M. Keynes,⁸ B. O. Koopman,⁹ I. J. Good,¹⁰ C. A. B. Smith,¹¹ F. Schick¹², H. E. Kyburg,¹³ and A. P. Dempster¹⁴ have pioneered in exploring alternatives to credal uniqueness as a condition on credal rationality. None of them, however, has rested content with a condition as weak as credal convexity. In any case, the ramifications of rejecting credal uniqueness remain to be explored.

As I shall explain later on in this chapter, my motivation for rejecting credal uniqueness as a condition on credal rationality is that it has objectionable philosophical ramifications for the revision of credal states. Substantially the same motivation argues in favor of replacing credal uniqueness by credal convexity, as I shall explain in chapter 9. It is now time to turn to the second of our four questions—the revision of credal states.

4.3

Confirmational Commitments

I have already noted that adoption of K as one's corpus of knowledge precludes endorsing some credal states which would be permitted were some other corpus adopted. Within the framework of restrictions on credal rationality I have imposed thus far, however, there is considerable room for choice among credal states relative to a fixed corpus of knowledge.

Carnap and Jeffreys envisaged constructing an inductive logic so powerful that all rational agents would be obliged to agree on the credal state to adopt relative to any given corpus of knowledge. I shall return to this point shortly. But even if the Carnap-Jeffreys program is rejected, we may suppose that rational X at time t is committed to a rule for choosing credal states relative to various potential corpora. Such a rule could be represented as a function $C(K) = B$, whose domain consists of the potential corpora and whose range consists of the credal

states. I shall call rules of this kind *confirmational commitments*.†

If we assume that rational X should adopt a confirmational commitment $C_{x,t}$, it follows from the intended interpretation that the following condition holds:

Total Knowledge Requirement: $C_{x,t}(K_{x,t}) = B_{x,t}$.

I know of no decisive argument in favor of requiring that rational X always be committed to a rule of the sort envisaged. However, the following considerations may be advanced in favor of such a view.

Let K' be the expansion of K obtained by adding e consistent with K to K and forming the deductive closure. Let B be a nonempty set of Q -functions satisfying the requirements of credal coherence relative to K and let B' be a nonempty set of Q' -functions satisfying credal coherence relative to K' . Then we can state that:

B' is the **Conditionalization of B with Respect to K and K'** if and only if for every $Q \in B$ there is a $Q' \in B'$ and for every $Q' \in B'$ there is a $Q \in B$ such that if f is consistent with K' , then $Q'(h; f) = Q(h; f \& e)$.

Suppose X at t endorses corpus $K_{x,t}$ and credal state $B_{x,t}$. Consider any potential corpus K' which is an expansion of $K_{x,t}$. Strict Bayesians maintain that, from X 's point of view at t , X should judge that if he were to shift from $K_{x,t}$ to K' by some system of transformations or other he should adopt as

† Confirmational commitments correspond approximately to credibilities in the sense of Carnap in "The Aim of Inductive Logic." However, a confirmational commitment need not be a "permanent disposition for forming beliefs on the basis of observations" (see p. 311 of "The Aim of Inductive Logic," in *Logic Methodology and Philosophy of Science*, edited by E. Nagel, P. Suppes, and A. Tarski, Stanford: Stanford University Press, 1962). I do not assume that X 's confirmational commitment is permanent. Furthermore, the body of evidence or knowledge which determines a credal state according to a confirmational commitment may contain theoretical assumptions, laws, or statistical assumptions, in addition to reports of observations. The agent is presumed to be committed to a rule for determining what his credal state (beliefs) should be for each potential corpus. The rule can be used to make revisions in credal state as long as it is held fixed. But not all revisions in corpus are expansions via observation and, hence, the rule provides for changes in credal state other than "on the basis of observations."

his credal state the conditionalization of $B_{x,t}$ with respect to $K_{x,t}$ and K' .

I shall offer an argument in favor of this requirement in chapter 10 based on decision-theoretic considerations. If the requirement is acceptable, it implies that once X has a credal state $B_{x,t}$ at t relative to his corpus $K_{x,t}$ he has a rule which specifies for each potential expansion of $K_{x,t}$ what his credal state should be.

To obtain the full rule defined for all potential corpora, we would have to extend the function thus defined to potential corpora which are not expansions of $K_{x,t}$. Such extensions are not determinable by reference to X 's current credal state and the conditionalizing arguments just noted without further supplementation.

However, in real life X will sometimes have views as to what his credal state should be relative to some contraction of his current corpus; and failures in X 's ability to identify a full range of commitments can be imputed to limitations on memory, computational facility, emotional stability, and other such disabilities.

I shall suppose that all rational agents adopt confirmational commitments and that such commitments satisfy the following condition:

Confirmational Conditionalization: If K' is a consistent expansion of K , $C(K')$ is the conditionalization of $C(K)$ with respect to K and K' .

Confirmational conditionalization is a condition on potential confirmational commitments. Adopting such a commitment at t does not obligate X to actually implement the dictates of the rule when he changes his corpus; but if X remains faithful to the confirmational commitment through such a change in corpus he should obey the rule.

Suppose X shifts from corpus K and credal state B at t to corpus K' and credal state B' , where K' is the expansion of K and B' is the conditionalization of B with respect to K and K' . Such a shift is a *temporal credal conditionalization*.

If X shifts from K' to K , where K is the contraction of K' (so that K' is the expansion of K) and B' is the conditionalization of B with respect to K and K' , a shift from B' to B is an *inverse temporal credal conditionalization*.

Confirmational conditionalization by itself does not entail that shifts in credal state due to expansion should be temporal credal conditionalizations, nor that shifts in credal state due to contraction should be inverse temporal credal conditionalizations.

Thus, anyone who insists that all revisions of credal state due to expansion should be temporal credal conditionalizations and all revisions due to contraction should be the inverse of this are committed to the view that rational X should remain faithful to the same confirmational commitment.†

Anyone who endorses this extreme view advocates *confirmational tenacity*. As shall become apparent, I reject confirmational tenacity.

Given confirmational conditionalization, a potential confirmational commitment can be defined by specifying the value of $C(UK)$. I shall call the probability measures relative to UK P -functions of the type $P(h; e)$. Since every potential K is an expansion of UK , if K is consistent, $C(K)$ is the conditionalization of $C(UK)$ with respect to UK and K .

As long as confirmational conditionalization is in place, therefore, the full force of credal coherence, consistency, and convexity can be secured by imposing these conditions ini-

† In *The Logic of Decision* (New York: McGraw Hill, 1965, pp. 160–161), R. C. Jeffrey argues that once a statement has probability 1, conditionalization precludes its probability being modified. If conditionalization is temporal credal conditionalization and the only sort of revision in credal state permitted is one conforming to temporal credal conditionalization, Jeffrey's claim is correct. But this is no argument for refusing to permit observation reports or any other extralogical statements from being assigned probability 1. This is good reason for abandoning temporal credal conditionalization as obligatory in all situations. Even if confirmational tenacity were mandated, X would not be obliged to conform to temporal credal conditionalization in all cases but only in those cases where his corpus is expanded. Jeffrey himself acknowledges this point on p. 154 of his book, where he writes that "there are cases in which a change in the probability assignment is clearly called for, but where the change is not occasioned simply by learning of the truth of some proposition E ." The most obvious example of this sort of situation emerges when the agent revises his corpus of evidence by contraction. Jeffrey, of course, was not thinking of this sort of change when he acknowledged that conditionalization (i.e., what I call temporal credal conditionalization) is restricted in applicability to cases of "learning of the truth of some proposition E "—i.e., to expansion. But that does not excuse him from, in effect, arguing against the feasibility of contraction by claiming that once a hypothesis is in a corpus and has probability 1 conditionalization precludes its removal. The circle should be obvious.

tially on $C(UK)$. Confirmational conditionalization guarantees that they will be satisfied by $C(K)$ for every potential K .

Since some authors (most notably Kyburg and Dempster)† reject confirmational conditionalization, it is important to keep in mind formulations of the several requirements on credal rationality that do not depend on confirmational conditionalization.

In any case, the various conditions on credal rationality already introduced and those to be considered subsequently may be understood to be constraints on potential confirmational commitments as well. In some contexts, it is helpful to think of them in that way.

4.4 Bayes' Theorem and Conditionalization

Let K and e (consistent with K) entail that at least and at most one of the hypotheses h_1, h_2, \dots, h_n is true and that each h_i is consistent with K and e . Let $Q \in B = C(K)$. The calculus of probabilities guarantees that if the Q -function is a probability measure as credal coherence requires, the following holds:

Bayes' Theorem: $Q(h_j; e) = Q(e; h_j)Q(h_j) / \sum_{i=1}^n Q(e; h_i)Q(h_i)$.

Strictly speaking, Bayes' theorem is a theorem of the calculus of probabilities. When the calculus of probabilities is applied to representation of Q -functions eligible for membership in credal states as credal coherence requires, Bayes' theorem becomes a consequence of credal coherence. In this last sense, Bayes' theorem is a condition on credal rationality. In neither sense does Bayes' theorem by itself provide any prescriptions for revising credal states.

Suppose, however, that K' is the expansion of K obtained by adding e , and let $B' = C(K')$. Confirmational conditionalization requires that B' be the conditionalization of $B = C(K)$ with respect to K and K' . Hence, for any $Q \in B$, there is a $Q' \in B'$ such that $Q'(h_j) = Q(h_j; e)$; and then, by Bayes' theorem,

$$Q'(h_j) = \frac{Q(e; h_j)Q(h_j)}{\sum_{i=1}^n Q(e; h_i)Q(h_i)}.$$

Thus, once confirmational conditionalization is imposed as a constraint on confirmational commitments, Bayes' theorem

† The ideas of Kyburg and Dempster will be considered in chapter 16.

can be employed to derive features of the Q -functions in the credal state relative to K' from features of the corresponding Q -functions in the credal state relative to K , where K' is an expansion of K .

Furthermore, if X remains faithful to his confirmational commitment during a shift from K to the expansion K' , temporal credal conditionalization will apply to the revision of credal state and Bayes' theorem can be used to compute the credal state relative to K' from the credal state relative to K .

Notice that the applicability of confirmational conditionalization and Bayes' theorem does not depend on the strict Bayesian demand that all credal states be representable by single Q -functions as credal uniqueness requires. Even if B relative to K contains several Q -functions, Bayes' theorem can be applied to every one of these functions to obtain Q -values for h_j conditional on e and, hence, a conditional Q -distribution over the h_i 's on the datum e . And confirmational conditionalization can then be used to determine a corresponding unconditional Q' -distribution over the h_i 's in the new credal state B' relative to K' .

Of course, strict Bayesians do insist on credal uniqueness; and the use of Bayes' theorem is customarily introduced with this requirement taken for granted. In order to lay the groundwork for later developments, it is important to understand that credal uniqueness is not presupposed by confirmational conditionalization, temporal credal conditionalization, its inverse, or the applications of Bayes' theorem in discussions of the revision of credal states.

Most self-styled Bayesians not only endorse credal uniqueness but take for granted that when X expands his corpus by adding datum e , the degree of credence to be assigned hypothesis h_j relative to the expanded corpus should equal $Q(h_j; e)$ where Q represents the initial credal state B and $Q(h_j; e)$ is computed from the values of $Q(h_i)$ for all the h_i 's and the values of $Q(e; h_i)$ for all the h_i 's via Bayes' theorem. For this reason, $Q(h_j; e)$ is called the *posterior probability* of h_j on e , $Q(h_j)$ the *prior probability*. $Q(e; h_j)$ is sometimes called the *likelihood* of h_j on e .

This terminology is somewhat misleading. Critics of Bayesian doctrine might be prepared to endorse credal coherence, and hence Bayes' theorem, as a condition on Q -functions

eligible for membership in credal states. Yet, they might reject confirmational conditionalization. In that case, Bayes' theorem may not play a fundamental role in an account of the revision of credal states. Talk of prior and posterior credal probability clearly becomes inappropriate.

I myself am willing to endorse confirmational conditionalization even though I reject confirmational tenacity. Consequently, I do not think that shifts in credal state due to expansion should be temporal credal conditionalizations in all cases. Yet, out of respect for tradition, I shall continue to talk of posterior and prior credal probability.

The chief point of interest here, however, is not a matter of terminological propriety, but the fact that there are many ways Bayesian doctrine on the revision of credal states can be modified without eviscerating it and leaving Bayes' theorem with no role to play.

4.5 Inductive Logic

The principles of inductive logic specify necessary conditions which a Q -function relative to K (or a P -function relative to UK) must satisfy in order to be eligible for membership in a credal state relative to K (relative to UK). As such they impose constraints on confirmational commitments. Moreover, they are context independent in that they are obligatory on all rational agents at all times and regardless of circumstances.

If a Q -function satisfies all the conditions imposed by a complete inductive logic IL relative to K , it is a *logically permissible* Q -function relative to K . A logically permissible P -function is, of course, a P -function satisfying the requirements of IL relative to UK .

Provided confirmational conditionalization is obeyed, the set of logically permissible P -functions determines the set of logically permissible Q -functions relative to potential K for every such K .

Jeffreys and Carnap entertained the ambition of developing principles of inductive logic IL so powerful that exactly one P -function would be logically permissible relative to UK (for given language L) and, hence, exactly one Q -function logically permissible relative to each potential corpus.¹⁵

The implications of this view are considerable. All rational agents would be obligated, insofar as they are able, to endorse

credal states which satisfy credal uniqueness and would be obligated to do so by principles of inductive logic.

Furthermore, all rational agents would be committed to adopting a single standard confirmational commitment (for the language L) by the principles of inductive logic. This commitment CIL would be the logical confirmational commitment representable by the set of all logically permissible P -functions and such that $CIL(K)$ consists of all logically permissible Q -functions relative to K . Because that set is single-membered for all potential K , we can represent the confirmational commitment CIL by the P -function which is the unique member $CIL(UK)$ and may with some propriety call that P -function a logical-probability function. Or we may follow Carnap's practice and regard it as a measure of degree of confirmation, understood as a measure of degree of logical probability.

Rational agents would also be obligated by inductive logic to obey confirmational tenacity. Given the applicability of confirmational conditionalization, all changes in credal state due to an expansion of corpus would have to be temporal credal conditionalizations and all changes in credal state due to contraction would have to be inverse temporal credal conditionalizations.

Consequently, if a complete inductive logic of the sort envisaged could be defended, all the main elements of the classical strict Bayesian view of credal rationality and the revision of credal states would be tenable.

I do not believe that an inductive logic as powerful as that envisaged by Carnap in his earlier writings on the subject can be defended. On this matter, my views conform to what is now the received opinion.

Some authors agree with De Finetti and Savage that a complete inductive logic is restricted to a principle of credal coherence. I shall call such a view *coherentist*.

Other authors would supplement credal coherence with a principle of *direct inference* which stipulates how knowledge of chances or objective probabilities determines credal judgments about the outcomes of trials on chance setups. An account of direct inference will be postponed until chapter 12 where it will be considered in connection with a discussion of chance or objective probability. However, many authors who seem to endorse the intelligibility of a conception of objective

probability or chance, such as R. A. Fisher,¹⁶ J. Neyman, E. S. Pearson,¹⁷ and H. Reichenbach,¹⁸ seem committed to some version of a principle of direct inference—as does Kyburg,¹⁹ who, while dispensing with notions of chance, introduces principles of direct inference from knowledge of relative frequencies in appropriately specified reference classes.

Objectivist inductive logic is restricted to credal coherence and direct inference. Objectivist inductive logic is insufficient for the purposes of the Jeffreys-Carnap program.

Other principles of inductive logic have also been advocated. I shall discuss some of them later. My own view is that an objectivist inductive logic is a complete inductive logic. In any case, few authors nowadays seriously maintain that a complete inductive logic allows exactly one P -function to be logically permissible relative to UK .

Suppose X is committed to corpus K at time t . Principles of inductive logic determine which Q -functions are logically permissible relative to K . If several Q -functions are logically permissible, then, as far as inductive logic is concerned, X is free to endorse a credal state consisting of any convex subset of them.

Credal consistency cuts down the degree of arbitrariness by obligating X to pick a nonempty subset. Credal uniqueness obligates X to pick a single-membered subset.

Thus, if we insist on remaining strict Bayesians and embracing credal uniqueness, X is obliged as a rational agent to choose one of several (typically infinitely many) numerically determinate credal states allowed by inductive logic relative to the corpus K .

The arbitrariness mandated in all of this may be reduced somewhat by requiring X to conform to the dictates of confirmational tenacity. That is to say, X should endorse a numerically precise confirmational commitment and remain loyal to it indefinitely. In this way, X is not compelled to make an arbitrary selection of a credal state relative to every shift in his corpus but only one grand arbitrary choice of a confirmational commitment.

I shall call anyone who embraces one of these views a *personalist* or an *intemperate personalist*. I am not sure whether intemperate personalism coincides precisely with the

views of any historical authors. However, it approximates views advanced by Savage and De Finetti.

One objection which might be raised against intemperate personalism is that X will often be incapable of identifying a numerical representation as accurately describing his credal state. Anxieties concerning this matter seem to have led many authors to explore ways and means of helping X elicit information about his credal state from data concerning his preferences, qualitative judgments of credal probability, and choices.

In my opinion, such problems are akin to worries about the implementability of the deductive closure requirement on X 's corpus of knowledge. They are, when rightly construed, of genuine importance. However, I understand strict Bayesians to be concerned with X 's commitments and not with the extent to which X succeeds or is capable at the moment of living up to his commitments. The same limitations on computational facility, memory, and emotional and social stability which pose obstacles to living up to the commitments imposed by deductive closure also pose obstacles to living up to the demands imposed by principles of credal rationality and the revision of credal states. I mean to bypass them here, as I did when discussing the standards for serious possibility.

My objection is to the prescriptions intemperate personalists insist on making as to what X 's commitments ought to be. Given his corpus K , X is in a situation, so intemperate personalists maintain, where he has no warrant for choosing one logically permissible Q -function rather than another as the sole member of his credal state. Yet, insofar as he is able, X should embrace a credal state containing exactly one such Q -function as its sole member.

Such arbitrariness would not be offensive, perhaps, if the choice of a Q -function made little difference in deliberation and inquiry. In point of fact, however, it will often make a considerable difference.

When X is in a predicament where he lacks a warrant for choosing one credal state rather than another, wisdom dictates that X should suspend judgment between the alternatives—i.e., select a credal state which expresses such suspension of judgment.

Thus, if the only considerations which warrant ruling out Q -functions are principles of inductive logic, X should not rule out any logically permissible Q -functions. He should endorse a credal state which avoids such elimination of Q -functions without warrant by suspending judgment between them.

A Q -function is *seriously permissible* according to X at t if and only if X at t has not ruled that Q -function out of his credal state $B_{X,t}$ relative to his corpus $K_{X,t}$. The set of Q -functions constituting a credal state is the set of Q -functions seriously permissible when X endorses that credal state as his own.

My thesis is that X 's credal state relative to K —i.e., the set of Q -functions seriously permissible according to X relative to K —should consist of all those Q -functions which X has no warrant for ruling out as impermissible relative to K . If the sole warrant for ruling out such Q -functions is that they fail to satisfy conditions of inductive logic, what I am recommending is that X 's credal state should consist of all Q -functions which are logically permissible relative to K .

There is yet another way to state this position. Recall that the logical confirmational commitment is the rule CIL which specifies for every K that $CIL(K)$ be the set of all Q -functions logically permissible relative to K . Because Carnap thought that inductive logic could be strengthened to the point where exactly one Q -function is logically permissible for every K , he held that all rational agents are obligated to endorse CIL.

But even if one abandons Carnap's view of inductive logic, one can still recommend that ideally rational agents endorse the weakest confirmational commitment CIL. One is not compelled to do so by inductive logic. However, two other assumptions provide a compelling basis for the requirement:

- (a) The contention that principles of inductive logic provide the sole warrant for ruling out Q -functions as not being seriously permissible.
- (b) The contention that one should not rule out Q -functions unless one has a warrant for doing so.

Savage called the Jeffreys-Carnap view that one should endorse a numerically definite CIL *necessary*.²⁰ I shall call any view which obligates rational X to embrace CIL *necessitarian*.

The import of necessitarianism differs depending on what one regards a complete inductive logic to be.

Thus, Keynes was a necessitarian; but he differed from Jeffreys and Carnap in that he refused to maintain that exactly one Q -function is logically permissible relative to potential K .

Keynes is regarded as at least partially within the Bayesian tradition. Fisher, Neyman, and Pearson are not. Yet, they recognize the propriety of using Bayes' theorem under certain conditions in deriving new appraisals of hypotheses with respect to what looks like credal probability from old appraisals prior to the acquisition of new data.

I favor interpreting these authors as being committed to an objectivist inductive logic. To the extent that credal coherence and direct inference fail to mandate a numerically precise Q -distribution over a set of alternative hypotheses, they all seem to suggest that no probability judgments are to be made. On the other hand, when direct inference and credal coherence can be used to make precise determinations, we are obliged to do so. I propose reading these authors as intending to say that when credal coherence and direct inference fail to mandate a numerically precise Q -distribution, all Q -distributions logically permissible according to an objectivist inductive logic ought to be seriously permissible.

Such a reading introduces some historical distortion. However, I do not think the distortion does serious injustice to the authors being considered and, if it does, it still remains the case that the views I am attributing to these authors are worth considering even if they did not hold them.

Necessitarianism avoids the excesses of intemperate personalism. Is it tenable? That, in large measure, depends on how powerful one can make inductive logic be without absurdity. If inductive logic is very powerful, the logical confirmational commitment will be fairly definite even if credal uniqueness is not strictly obeyed. This may turn out to suffice for the uses of credal states in decision theory and scientific inquiry. Furthermore, necessitarianism entails confirmational tenacity and, hence, provides an account of how credal states should change with changes in knowledge. An account of the revision of credal states reduces to an account of the revision of corpora of knowledge.

If, however, inductive logic is very weak, CIL may not be

able to be used to furnish credal states helpful in deliberation and inquiry and, worse yet, there may be little hope that through further inquiry one could obtain a credal state which would be helpful. Necessitarianism would counsel a sort of skepticism which would render it unfeasible to carry on deliberation and inquiry.

I adopt an objectivist inductive logic. Chapters 15-17 argue that objectivist necessitarianism is untenable and that variant forms of necessitarianism are no better.

Must we then reject thesis (b) and embrace an intemperate personalism? In my opinion, we should explore ways and means of modifying condition (a) so that (b) can be obeyed. This alternative will now be considered.

4.6 Contextualism

Intemperate personalists and necessitarians share an important assumption in common. It is contention (a) of section 4.5, which asserts that the only justification for ruling out a Q -function as not being seriously permissible relative to one's corpus K is that it is not logically permissible relative to K .

Intemperate personalists and necessitarians do disagree concerning contention (b). Intemperate personalists require rational agents to choose a single Q -function as the sole seriously permissible one relative to K without warrant. Necessitarians insist that no Q -function should be ruled out unless there is some warrant for doing so.

The principles of inductive logic, it may be recalled, impose constraints on Q -functions eligible for membership in credal states relative to potential corpora obligatory on all agents at all times. They are context-independent principles in the sense that ideally rational X is committed, insofar as he is able, to endorsing a confirmational commitment conforming to the requirements of inductive logic regardless of what his corpus of knowledge, his goals and values, or other circumstances happen to be.

To be sure, principles of inductive logic fix credal states in a context-dependent manner to the extent that the logical permissibility of a Q -function is relative to a potential corpus. But confirmational commitments are controlled by such principles in a context-independent way.

Suppose, however, that the confirmational commitment X should endorse at t' depends on the following factors:

- (i) The corpus $K_{X,t}$ that X endorses at the prior stage t .
- (ii) The confirmational commitment $C_{X,t}$ that X endorses at t (which together with $K_{X,t}$ determines $B_{X,t}$).
- (iii) X 's goals and values, the problems he is investigating, the way he has succeeded in identifying potential solutions, and other circumstantial factors.

A contextualist might adopt the view that $C_{X,t'}$ should remain the same as $C_{X,t}$, everything else being equal—i.e., unless there is some relevant feature in the vaguely specified factors of type (iii) that warrants a change. Of course, if one is going to take contextualism seriously, one would have to state more precisely what factors of type (iii) are relevant and how they direct revisions of confirmational commitments from t to t' given $K_{X,t}$ and $C_{X,t}$.

However, we do not have to spell out details to see that if such a contextualist view is endorsed, contention (a) of section 4.5 will have to be abandoned. The Q -functions which are ruled out as not being seriously permissible at t' may be ruled out with justification, even though the justification relies on more than principles of inductive logic. Appeal may be made to the prior confirmational commitment and corpus and to factors of the type specified vaguely under (iii).

In his admirable "Scientific Inference," A. Shimony has outlined an approach to probabilistic judgment based on an explicitly contextualist outlook.²¹ Shimony acknowledges, however, that supplementing principles of inductive logic by reference to the problems under consideration and potential solutions identified will, in general, fail to single out a numerically precise confirmational commitment or even a numerically precise credal state relative to the corpus at the new stage t' .

In chapter 13, I shall outline a view of how contextual considerations control the adoption of confirmational commitments and their revision. Although this proposal differs in technical details from Shimony's and is more closely tied to issues of inferential expansion than his is, in a rough and ready way our views concerning the sorts of contextual considerations that should be considered relevant seem to be quite similar.

However, Shimony and I disagree concerning how to respond to the fact that contextual considerations will fail, in general, to single out a unique Q -function relative to K as eligible for consideration as seriously permissible. Shimony follows the intemperate personalists in insisting that rational X should obey credal uniqueness and adopt a credal state allowing exactly one Q -function to be seriously permissible even though it is not justifiably distinguished from the other Q -functions which survive criticism from inductive logic and contextual considerations. Shimony calls himself a *tempered personalist*. As I understand him, this means that unlike intemperate personalists he does take contextual considerations into account.

I, for my part, sympathize with the necessitarian insistence that contention (b) of section 4.5 be upheld. Although I side with Shimony against both intemperate personalists and necessitarians in regarding appeals to context as providing relevant warrants for retaining or revising confirmational commitments, I side against both intemperate and tempered personalists and with necessitarians in insisting that rational X should refuse to rule out Q -functions relative to K which survive criticism from legitimate considerations. I shall call such a view *revisionist*.

In section 2.7, I took the position that when two or more expansion strategies are admissible according to the principle of maximizing expected utility, one should avoid arbitrariness by adopting the strategy that allows for suspension of judgment between the admissible strategies. In section 3.3, I considered how contraction strategies should be evaluated. There I took the view that one should seek to minimize the loss of informational value which must be incurred, but that where two or more strategies minimize the loss one should implement both or all of them. In this way, one expresses one's suspension of judgment between these alternatives.

Revision of confirmational commitments is not quite the same sort of problem as revision of corpora of knowledge. Yet, in all our activities, one should avoid arbitrary decisions insofar as this is feasible. This core idea lurks behind my advocacy of contention (b) of section 4.5 and my concern to reject both intemperate and tempered personalism.

4.7
Rational
Valuation

Suppose X takes for granted in his corpus of knowledge that A_1, A_2, \dots, A_m are options feasible for X but that X is constrained to implement exactly one of them by some time t' .

For each feasible option A_i , let $c_{i1}, c_{i2}, \dots, c_{im_i}$ be a set of hypotheses such that the expansion K^* obtained by adding " A_i is implemented" entails that at least and at most one of the c_{ij} 's is true and such that each c_{ij} is consistent with K^* and, hence, the truth of "If A_i is chosen, c_{ij} is true" is a serious possibility according to X .

X 's goals and values induce a system of valuations of the c_{ij} 's representable by a set $G_{N,t}$ of real-valued functions $u(c_{ij})$ defined over all the c_{ij} 's. $G_{N,t}$ should satisfy the following requirements:

Valuational Consistency: $G_{N,t} \neq \emptyset$.

Closure under Linear Transformation: If $u \in G_{N,t}$ and if $u' = \alpha u + \beta$ for every positive real value α and for every real β , then $u' \in G_{N,t}$.

Valuational Convexity: If u and u' are in $G_{N,t}$ then so is every weighted average of u and u' (with positive weights).

All three requirements are endorsed by strict Bayesians. However, strict Bayesians would insist on the following strengthening of Valuational Convexity:

Valuational Uniqueness: If u and u' are in $G_{N,t}$, then there is a positive real α and a real β such that $u' = \alpha u + \beta$.

Strict Bayesians may acknowledge that even rational agents may lack the self-awareness to be able to identify their values precisely. But this sort of failure is analogous to the sort we have been neglecting all along. What strict Bayesians do seem to claim, however, is that rational agents ought to be so decisive in their values as to be committed to a system of values representable by a utility function unique up to a linear transformation.

When X is some institution or group, it is, of course, widely acknowledged that this view is problematic. Collective decision making is a case where the group is often in conflict with regard to its values. Many different proposals have been explored for identifying some ranking of alternatives, or even a utility function unique up to a linear transformation, that rep-

resents the group resolution of the conflict, on the basis of which decisions should be taken.

However, it is far from clear that collective agents are irrational if they fail to rank outcomes so as to remove all conflict or to represent their group valuation by means of a utility function unique up to a linear transformation.

What is true of collective agencies is applicable to persons as well. Men often face decision problems under circumstances where their values and goals are in conflict and there is no opportunity to find a basis for resolving the conflict prior to making a decision.

In such cases, I contend that it is wrong to impose as a condition of rationality that agents (whether they are persons or communities) should arbitrarily pick a resolution of conflict in value as the basis for decision. As in the case of credal states, one should not reach unwarranted and arbitrary conclusions. It is preferable to remain in suspense.

I interpret a system of values representable by a set of utility functions unique up to a linear transformation to be one where all value conflicts are resolved. If there is unresolved conflict, $G_{N,t}$ should contain two or more subsets of u -functions each of which satisfies strict Bayesian requirements.

Finally, if there is unresolved conflict, I contend that X should avoid ruling out any potential resolution of the conflict between u -functions which have not been ruled out. By reasoning which shall be explained in chapter 8, this informal requirement supports the convexity requirement on $G_{N,t}$.

4.8
Maximizing
Expected
Utility

Given a decision problem of the sort specified in section 4.7, the expected utility of a feasible option A_i relative to a Q -function in $B_{N,t}$ and a u -function in $G_{N,t}$ is given by

$$E(A_i; Q, u) = \sum_{j=1}^{n_i} Q(c_{ij}; A_i \text{ is implemented})u(c_{ij}).$$

The method of *ranking options with respect to expected utility (REU)* stipulates that feasible option A_i outranks A_i' with respect to expected utility relative to Q and u if and only if $E(A_i; Q, u) > E(A_i'; Q, u)$ and ranks with A_i' if and only if $E(A_i; Q, u) = E(A_i'; Q, u)$.

In chapter 5, alternative methods for ranking feasible options with respect to expected utility relative to Q and u will

be considered. The following definitions are intended to apply regardless of which method for ranking feasible options is employed from among those considered.

A_i is *optimal* with respect to expected utility relative to Q and u if and only if there is no other feasible option which outranks A_i with respect to expected utility relative to Q and u .

A_i is *E-admissible* relative to $B_{N,t}$, $K_{N,t}$, and $G_{N,t}$ if and only if there is a seriously permissible Q -function in $B_{N,t}$ and a seriously permissible u -function in $G_{N,t}$ such that A_i is optimal with respect to expected utility relative to that Q -function and u -function.

An option is *admissible* if and only if rational X is entitled as a rational agent to choose it from among the available options given his corpus, credal state, and goals. A principle or criterion of rational choice specifies necessary and sufficient conditions for such admissibility. One such criterion would be the following:

The Principle of E-Admissibility: All and only those feasible options which are *E-admissible* are admissible.

If $B_{N,t}$ and $G_{N,t}$ obey conditions imposed by strict Bayesians, the principle of *E-admissibility* reduces to the following:

The Principle of Maximizing Expected Utility: All and only those feasible options which bear maximum expected utility relative to the uniquely permissible Q -function and uniquely permissible u -function are admissible.

The principle of maximizing expected utility becomes inapplicable when strict Bayesian conditions on credal states and goals and values are modified. The principle of *E-admissibility* is a natural generalization of the principle of maximizing expected utility which covers all the new cases. It has been suggested by Good.²²

In my opinion, neither the more general criterion nor its specialization are acceptable.

In some decision problems, more than one feasible option may be *E-admissible*. This can happen even when strict Bayesian conditions are satisfied. Presumably X is compelled to choose some feasible option. Hence, as far as considerations

of *E-admissibility* are concerned, there is no basis for rendering a verdict.

It appears, therefore, that at the level of choice, an element of arbitrariness appears which cannot be avoided by suspending judgment.

However, there are, in my view, some sorts of decision problems where some options feasible for the agent express suspension of judgment. For example, if X adds $h \vee g$ to his corpus K together with the deductive consequences, he suspends judgment between h and g and, in that sense, his expansion strategy expresses suspension of judgment between adding h to K with all the deductive consequences and adding g to K with all the deductive consequences.

There are other contexts where implementing one option may express suspension of judgment between others. This is true in contracting a corpus of knowledge. It may also be true in some contexts of practical decision making as well. However, it may be one of the salient differences between cognitive decision making and various types of practical decision making that in the former sort of decision making, suspension of judgment between feasible options is expressible by another feasible option.

I contend that when two or more feasible options are *E-admissible* and there is a feasible option which expresses suspension of judgment between them, it should, everything else being equal, be chosen.

A full articulation of this principle will be postponed to chapter 6. However, endorsing it entails rejecting the principle of *E-admissibility*.

I have already rejected that principle in the manner indicated in introducing a rule for ties to arbitrate between rival *E-admissible* expansion strategies (see section 2.7). In chapter 6, the full ramifications of replacing the principle of *E-admissibility* by an alternative in the context of cognitive decision making will be discussed.

A full characterization of the alternative rule must be deferred. The gist of the proposal is this: The notion of a *P-admissible* option is introduced. If there is a feasible option which expresses suspension of judgment between all *E-admissible* options and it satisfies certain other requirements, it is

P-admissible. If there is no such option, all *E*-admissible options are *P*-admissible.

An *S*-admissible option is an option which is maximin among all the *P*-admissible options. A full discussion of this idea will be delayed to chapter 7. It is designed primarily to cover cases where all *E*-admissible options are *P*-admissible.

Decision theorists are accustomed to distinguish between decision problems under risk where credal states are strictly Bayesian (as are goals and values) and decision problems under uncertainty where, in effect, credal states are maximally indeterminate—i.e., where all *Q*-functions obeying credal coherence are permissible.

Several criteria for rational choice have been proposed for decision making under uncertainty. A favorite is maximin. I myself am inclined to favor a suitably modified version of that principle for decision making under uncertainty.

However, my main thesis is that whatever principle we use for decision making under uncertainty ought to be so understood that it applies to situations intermediate between decision making under risk and decision making under uncertainty.

Hence, I take an *S*-admissible option to be a maximin solution from among all *P*-admissible options rather than all feasible options. In the case of decision making under complete uncertainty, all “undominated” feasible options are *E*-admissible and *P*-admissible so that *S*-admissible options according to my proposal become maximin solutions in the usual sense.

Thus, if we endorse the condition that *S*-admissibility is necessary for admissibility, we manage to provide an account of rational decision making which respects the desirability of avoiding arbitrary choice wherever feasible, conforms otherwise to the dictates of strict Bayesianism in one specifiable class of cases, conforms to a familiar theory of decision making under uncertainty in another class of special cases, and covers all sorts of intermediate cases as well.

4.9 Revision

The chief attraction of strict Bayesian doctrine has been its capacity to provide a systematic approach to the four questions posed at the beginning of this chapter. As long as no alternative is offered which meets objections to strict Bayesianism while posing an alternative unified approach to these

questions, strict Bayesians will continue to have the excuse that nothing better is available. One of my aims in this chapter has been to show, at least in outline, that there is an alternative which offers systematic answers to the four questions while meeting some of the difficulties confronting strict Bayesianism.

Bayesianism, however, ought not to be given more credit than it deserves. There is no decisive strict Bayesian answer to all aspects of question (b) of revising cognitive resources. Even Jeffreys and Carnap pretended to do no more than provide an account of how *X*'s credal state should be determined by his corpus of knowledge and a confirmational commitment mandated by very powerful principles of inductive logic. No account is offered of how a corpus of knowledge should be revised.

To be sure, Carnap sometimes leaves the impression that he is committed to the view that one may revise a corpus by routine expansion via observation and, perhaps, by some sort of conceptual change.²³

Notice, however, that these epistemological views of his are a supplement of his strict Bayesian doctrine, and do not follow from it. There are alternative approaches to the revision of corpora which could be advanced which do not conflict in any critical way with the Carnap-Jeffreys approach to the revision of credal states, given changes in corpus. One could provide for inferential or inductive expansion as well as routine expansion. One could provide for contraction.

What is crucial, however, is that a comprehensive approach to question (b) does require some account of the revision of corpora of knowledge. No strict Bayesian approach to revision of credal states, whether it is necessitarian like Carnap's, intemperately personalistic like De Finetti's, or of the tempered personalist variety advocated by Shimony, can avoid acknowledging that as long as there are some contexts where the legitimacy of revisions of credal states depends on obedience to confirmational tenacity and conformity with confirmational conditionalization, that legitimacy also depends on the propriety of making some sort of change in corpus.

To be sure, those strict Bayesians who abandon the Carnap-Jeffreys project are no longer obliged to insist on conformity to confirmational tenacity on all occasions—although most

personalists seem to think that as a rule one should be tenacious in keeping one's confirmational commitment fixed.

But this only compounds the problem; for now the revision of credal states depends for its legitimacy on two factors: (i) whether the corpus is legitimately revised or not and, if so, how; and (ii) whether the confirmational commitment is legitimately revised or not and, if so, how.

Even less attention is devoted to the problem of revising confirmational commitments than to the problem of revising corpora of knowledge.

One cannot object to strict Bayesians who disclaim any intention of investigating the twin problems of revising corpora of knowledge and confirmational commitments. One does not have to answer every question in order to make headway on some.

The impression ought not to be left, however, that Bayesians do qua Bayesians provide complete answers to our four questions. Even the answers offered to the question of revising credal states which depend on the use of conditionalization and Bayes' theorem are only partial and incomplete.

In this chapter, I have tried to explain why answers to the problem of revising confirmational commitments meeting strict Bayesian requirements are unsatisfactory from my point of view. That outline requires some elaboration and a positive approach to the revision of confirmational commitments should be proposed. I shall attempt this in chapter 13.

My approach to the revision of confirmational commitments involves the account I favor of inferential expansion. I have already discussed this approach in chapter 2. However, that discussion presupposes that strict Bayesian requirements on credal states and valuations in terms of epistemic utilities are satisfied. In chapter 6, I shall elaborate upon the outlook of chapter 2 by explaining the ramifications for inferential expansion of abandoning strict Bayesian requirements.

4.10 Prospectus

The remaining portions of this book are devoted to elaborating on and defending a revisionist outlook towards changing credal states. There are several loose ends to the argument which should be picked up, however, before I attempt to offer an account of the revision of confirmational commitments.

In chapter 5, I consider some relatively technical questions

concerning the conditions which ought to be imposed on probability measures in the condition of credal coherence. Should a logically permissible Q -function relative to corpus K always assign positive value to a hypothesis consistent with K ? Should $Q(h; e)$ be defined only for e bearing positive Q -value? Should countable additivity be imposed as a requirement on probability measures in the condition of credal coherence? I answer all these questions in the negative. Chapter 5 explains the motives, outlines methods for handling some technical complications, and discusses some modifications which are required in the definition of E -admissibility because of the position I take.

Chapter 6 elaborates on the notion of suspending judgment between options and explicates the concept of P -admissibility. The results are applied to inferential expansion.

Chapter 7 explicates S -admissibility and the role of maximin in the decision theory being proposed.

Chapter 7 concludes the discussion of the general principles of the theory of rational choice I am proposing. No effort will be made to prove that the fundamental principles underlying that theory are correct. However, given the overall theory, some further explanation of how the conditions imposed on rational valuation and, in particular, valuational consistency and convexity are desirable components of the theory may prove helpful. This explanation is offered in chapter 8 in the context of a formal elaboration of the concept of conflict between values.

Chapter 9 focuses on the rationale of the requirement of credal convexity and some of its ramifications.

Chapter 10 offers a rationale for confirmational conditionalization, relying on the elements of the decision theory already adopted.

I have already stated that I favor an objectivist inductive logic consisting of a principle of direct inference and the principle of credal coherence. To formulate the principle of direct inference I favor, however, it is first necessary to discuss the concept of chance (also called statistical or objective probability). Chapter 11 discusses preliminary notions of ability and disposition presupposed by my account of chance. Chapter 12 introduces the conception of chance and with it the principle of direct inference I favor.

Chapter 13 introduces an account of the revision of confirmational commitments based on the assumption that objectivist inductive logic is complete. This chapter concludes the discussion of the positive proposals I have to make concerning the revision of credal states and confirmational commitments.

In my opinion, the important alternatives to objectivist revisionism are various forms of necessitarianism and, in particular, objectivist necessitarianism. If an acceptable account of deliberation and inquiry could be constructed on an objectivist necessitarian foundation, there would be no need to furnish an account of the revision of confirmational commitments based on an appeal to contextual considerations. All rational agents would be obliged to endorse the same standard confirmational commitment.

The first two sections of chapter 13 explain some serious objections to objectivist necessitarianism including the charge that the testimony of the senses cannot be used effectively by objectivist necessitarians in deliberation and inquiry.

One way to meet the difficulty while remaining necessitarian is to strengthen inductive logic. Chapters 14 and 15 explore ways of doing so through rationalizing fiducial inference—in particular, by introducing I. Hacking's law of likelihood. Utilizing results obtained by T. Seidenfeld, Hacking's proposal and others resembling it will be shown to lead to inconsistency.

An alternative approach is to retain an objectivist inductive logic but modify confirmational conditionalization. This idea is implicit in Fisher's own outlook on fiducial inference. Kyburg offers a modification of the principle of direct inference which entails deviation from confirmational conditionalization. A variant approach leaves direct inference unmodified but alters confirmational conditionalization. Such a view is suggested by proposals made by Dempster. Chapter 16 discusses Kyburg's view and the reconstruction of Dempster's view just mentioned.

Chapter 17 discusses a version of the Neyman-Pearson-Wald approach to statistical theory which remains loyal to the strict letter of objectivist necessitarianism, accepts the allegedly offensive consequences of doing so, but denies that these consequences imply the uselessness of the testimony of the senses. Ways and means are devised to circumvent the charge

that data become useless by providing a new role for the testimony of the senses in inquiry and deliberation alternative to the role it plays according to those who emphasize the importance of using data with Bayes' theorem and conditionalization to modify judgments of credal probability.

Chapters 13–17 are designed to serve two purposes:

First, the diverse responses to the challenge to objectivist necessitarianism posed in chapter 13 are shown to be inadequate, thereby undermining this alternative to objectivist revisionism.

Second, whether the objectivist revisionist view I favor or some alternative is endorsed, the method of representing credal states by convex sets of probability distributions employed in this discussion affords an opportunity to discuss a wide variety of different and often conflicting positions on probability judgment in a manner which identifies shared assumptions and points of disagreement within a single framework. Anyone who has examined the topics discussed in this essay will appreciate how difficult it is to discuss diverse viewpoints without seeming to beg questions in one way or another.

I do not pretend to avoid begging all questions. Yet, the framework for discussing probability judgment developed here can be employed by those who defend substantially different views from my own to develop a common discourse within which some progress may be made in clarifying some of the issues under dispute. Even those who remain unconvinced by my conclusions might find this feature of the discussion helpful.

**5.1
Credence and
Choice**

Credal states function together with standards for serious possibility as resources in deliberation and inquiry. They are used chiefly to evaluate feasible options with respect to *E*-admissibility. They serve in this capacity in all contexts of deliberate decision making, both "practical" and "cognitive," where the concern is with the revision of cognitive states.

The several requirements imposed on credal states as conditions of credal rationality are, to a large degree, a reflection of the use of *Q*-functions in credal states in computing expected utilities and, hence, in determining *E*-admissibility.

In this chapter, a brief accounting of the relations between *E*-admissibility and the conditions of credal coherence and consistency will be offered. The bulk of the discussion will be devoted to explaining why these relations do not warrant strengthening the condition of credal coherence.

I shall first explain why none of the following requirements are appropriate: (1) that *Q*-functions in credal states be regular probability measures, (2) that for every *Q*-function, $Q(h; e)$ be defined only for *e* such that $Q(e) > 0$, and (3) that *Q*-functions be countably additive measures.

Persuasive arguments have been advanced for insisting upon a regularity condition derived from considerations of the role of *Q*-functions in computing expectations of the sort I shall invoke in defense of credal coherence. I shall show that these arguments have relied, in fact, on an oversight in discussions of how feasible options are to be ranked with respect to expected utility; when that oversight is corrected, these arguments lose their force.

Finally, I shall explore ways of representing *Q*-functions that violate countable additivity, which may prove helpful in clarifying some recent controversies concerning "improper" distributions. The account given in this chapter is preliminary

to further discussions which will be found in sections 12.14, 12.16, and 15.9.

Worries about countable additivity are provoked by technical concerns with infinity. Those readers who are not anxious may skip sections 5.9–5.11, returning to them only if some issue in the subsequent discussion should suggest a need to do so.

**5.2
Classical Rep-
resentations of
Decision Prob-
lems**

The normative account of rational choice proposed here is intended to apply to contexts of deliberate and rational choice where agent *X* knows he is compelled to implement at least and at most one option from set \mathcal{A} but is free to choose which one it will be. Each $A_i \in \mathcal{A}$ is associated with a set \mathcal{C}_i of hypotheses. *X* knows that if he implements A_i , at least and at most one hypothesis from \mathcal{C}_i is true. Moreover, for every $c_{ij} \in \mathcal{C}_i$, the statement "if A_i is implemented then c_{ij} is true" is consistent with what *X* knows.

Usually decision problems are formulated somewhat differently. Given the set \mathcal{A} , and a set *U* of hypotheses exclusive and exhaustive relative to *K* and each consistent with *K* (the states of nature), for each $A_i \in \mathcal{A}$ and each $h_k \in U$, there is an hypothesis o_{ik} such that *K* and the information that A_i is implemented entails that h_k is true if and only if o_{ik} is true. For given A_i , the set \mathcal{O}_i then serves as the set \mathcal{C}_i of hypotheses concerning the consequences of implementing A_i .

Representations of decision problems of this kind shall be called *classical* representations. Given a representation of a decision problem of the first sort, it is possible to obtain a classical representation.

I am assuming that \mathcal{A} and the \mathcal{C}_i 's are all finite and that, therefore, there is a largest finite set \mathcal{C}_i^* bearing the number n_i^* .

Consider any other option $A_{i'}$ and $\mathcal{C}_{i'}$ with $n_{i'} < n_i^*$. One can take any element $c_{i'j}$ of $\mathcal{C}_{i'}$ and refine it into a disjunction of hypotheses $h_1 \& c_{i'j} \vee \dots \vee h_r \& c_{i'j}$, equivalent given *K* and " $A_{i'}$ is implemented" to $c_{i'j}$. In this way, we can always extend the set $\mathcal{C}_{i'}$ so that its cardinality becomes n_i^* . Moreover, we should always be able to do this so that for every permissible *u*-function in *G*, the disjuncts all bear the same *u*-value as $c_{i'j}$. Having done this for each A_i , we have for each

feasible option a system of hypotheses about consequences, each system bearing the same number of elements. Enumerate them so that for each feasible option there is a first, second, etc. hypothesis. Let U consist of n_i^* exclusive and exhaustive alternatives such that h_k asserts that for some i , A_i is implemented and the k th consequence for A_i is true.

There are, of course, many ways of carrying out the instructions of this recipe. Moreover, there will sometimes be other ways of obtaining classical representations as well. There will be no problem in using classical formulations and no dearth of them for a given decision problem. I shall, for the most part, discuss decision problems under the assumption that they are given a classical formulation.

Given any u -function in G and any Q -function in B , then

$$E(A_i; Q, u) = \sum_{k=1}^n Q(o_{ik}; A_i \text{ is implemented})u(o_{ik})$$

for any decision problem with a classical formulation. Since

$$Q(o_{ik}; A_i \text{ is implemented}) = Q(h_k; A_i \text{ is implemented}),$$

then

$$E(A_i; Q, u) = \sum_{k=1}^n Q(h_k; A_i \text{ is implemented})u(o_{ik}).$$

This formula may be substituted for the one introduced in section 4.8 for the purpose of applying REU to rank feasible options with respect to expected utility. The earlier formula may be used regardless of whether the problem is formulated classically or not. This one presupposes a classical representation of the decision problem.

5.3 Dominance

Relative to a classical formulation of a decision problem with states of nature in U , $A_i \in \mathcal{A}$ *strongly dominates* $A_{i'} \in \mathcal{A}$ if and only if for every $u \in G$ and every $h_k \in U$, $u(o_{ik}) = u_{ik} > u_{i'k} = u(o_{i'k})$.

Relative to a classical formulation, A_i *weakly dominates* $A_{i'}$ relative to G and U if and only if $u_{ik} \geq u_{i'k}$ for every $h_k \in U$, and $u_{ik} > u_{i'k}$ for some $h_k \in U$.

Thanks to R. C. Jeffrey and R. Nozick,¹ it is now widely recognized that there are circumstances when strongly dominated options may be admissible while there are others where they are not. Indeed, that should be obvious; for options may

be strongly dominated relative to one classical formulation of a decision problem and fail to be dominated relative to another.

Some controversy exists concerning the necessary and sufficient conditions that a classical formulation should satisfy for a strongly dominated option to be forbidden as not E -admissible.

There is no need to engage in the controversy here. Identification of a condition which is widely acknowledged to be sufficient for enjoining against allowing strongly dominated options to be E -admissible will help explain why Q -functions used to compute expected utilities should be required to conform to the requirements of the calculus of probabilities.

The states of nature in U and feasible options in \mathcal{A} are Q -independent of one another relative to a given Q -function if and only if $Q(h_k; A_i \text{ is implemented}) = Q(h_k)$ for every $A_i \in \mathcal{A}$ and every $h_k \in U$.

When Q -independence holds relative to Q , then

$$E(A_i; Q, u) = \sum_{k=1}^n Q(h_k)u(o_{ik}) \text{ for every } u \in B.$$

The options in \mathcal{A} are *credally irrelevant* to the states in U if and only if all states in U are Q -independent of the options in \mathcal{A} relative to each permissible Q -function in B .

The prohibition against choosing a strongly dominated option is widely thought to be legitimate when the feasible options are credally irrelevant to the states of nature.

To guarantee that in this class of cases no strongly dominated option is admissible, the following stronger condition will be adopted:

The Principle of Strong Dominance (SD): In a classical formulation of a decision problem where states in U are credally independent of the options in \mathcal{A} , no option strongly dominated by another in \mathcal{A} is E -admissible.

The conditions on P -admissibility and S -admissibility to be explained in later chapters will guarantee that conformity with the principle of strong dominance secures that no strongly dominated option will be admissible.

5.4
Credal
Coherence

Consider any decision problem where the options are credally irrelevant to the states and modify the problem by selecting any Q -function from the credal state and letting the new credal state be the set whose sole member is that Q -function. The new decision problem must also satisfy the condition of credal irrelevance.

In the same spirit, modify the set G by shifting to a new set consisting of linear transformations of exactly one u -function in G .

The result is a strictly Bayesian decision problem which, nonetheless, is regulated by SD .

Finally, any option $A_{i'}$ strongly dominated by A_i in the original decision problem will remain so in the modified strictly Bayesian decision problem no matter which Q -function and u -function is picked.

In the strictly Bayesian decision problem so obtained, all E -admissible options will be optimal with respect to Q and u where optimality is determined by the ranking with respect to expected utility relative to Q and u .

For the present, we have been assuming that REU regulates this ranking.

From these considerations, it follows that if A_i strongly dominates $A_{i'}$ in the original problem and, hence, in the modified problem, A_i outranks $A_{i'}$ relative to Q and u . This is due to SD .

REU entails that

$$E(A_i; Q, u) > E(A_{i'}; Q, u).$$

Hence,

$$\sum_{k=1}^n Q(h_k)u_{ik} > \sum_{k=1}^n Q(h_k)u_{i'k}.$$

Thus, Q -values assigned the h_k 's must satisfy conditions implied by this inequality for every u -function we might envisage where the strong dominance of $A_{i'}$ by A_i obtains.

If we assume that condition (2) on Q -functions (section 4.2) which are probabilities should be satisfied to secure univocality, the unconditional Q -distribution over elements of U can be shown to take nonnegative values as (1) requires, finite additivity as in (3) is satisfied and the disjunction of all elements of U must take some positive and finite Q -value a so

that dividing the Q -values for hypotheses equivalent given K to Boolean combinations of elements of U by a yields a new function which not only satisfies (1), (2), and (3) for unconditional Q -functions, but (4) as well.

These results are due essentially to F. P. Ramsey and De Finetti.² However, they formulated their arguments in a manner which entailed that credal uniqueness is a condition on credal rationality. I have done no such thing.

Stripping away their tacit assumption of credal uniqueness, we see that they showed that SD and REU mandate that Q -distributions over the states which are unconditional or conditional on the options conform to conditions (1)-(4) of section 4.2.

It also possible to show that, given REU and credal coherence (or, at any rate, that part of credal coherence which implies that when credal irrelevance of states on acts obtains, the unconditional Q -distributions over the states obey conditions (1)-(4)), SD obtains.

In this sense, a case can be made out for at least part of the condition of credal coherence on the grounds that it secures that SD is satisfied when credal states are characterized as functioning in inquiry and deliberation as furnishing "weights" for computing expected utility needed in applying REU to assess E -admissibility. It must be emphasized that this argument does presuppose that credal states do have a certain kind of function in deliberation and inquiry. Anyone who withholds this function from credal states or questions SD is free to reject the argument.

5.5
Weak Domi-
nance and
Credal Regu-
larity

In the case of classical representations of decision problems where options are credally irrelevant to hypotheses in U , we could envisage forbidding the admissibility of options which are weakly dominated by other feasible options and by an argument parallel to that given previously introduce the following requirement:

Weak Dominance (WD): In a classical formulation of a decision problem where options in \mathcal{A} are credally irrelevant for states in U , no option weakly dominated by another in \mathcal{A} is E -admissible.

Should weak dominance be obeyed? I think so. Surely of two options, one of which is sometimes (i.e., for some serious possibilities) better than the other and is never worse, that option ought to be recommended. Yet, as Shimony showed,³ conformity to REU and *WD* entails that unconditional *Q*-distributions over *U* and Boolean combinations thereof should assign positive values to all such hypotheses consistent with *K* as the conditions of credal regularity mentioned in section 4.2 requires. Shimony concludes that credal regularity should be imposed on *Q*-functions in credal states; Carnap seconds this view.⁴

De Finetti has explicitly rejected *WD* on the grounds that it leads to a commitment to credal regularity which he rejects.⁵

I think De Finetti is right in rejecting credal regularity as a condition which the *Q*-functions in credal states of rational agents must satisfy. Consideration of the following cases should help explain why.

Case 1: Suppose *X* considers the problem of estimating the value of a real-valued parameter which ranges between k_* and k^* . He might, for example, be concerned to estimate the unknown chance of obtaining heads on a toss of a coin. (This is not an example which De Finetti would use.) *X* does not know which of the noncountably many possible values is the true one. Any hypothesis specifying that such a value is the true one is a serious possibility according to him.

If credal regularity were to obtain, at most countably many of these hypotheses could bear positive *Q*-value if the other requirements on credal coherence are to be obeyed. Thus, either it must be admitted that some seriously possible hypotheses bear 0 *Q*-value or situations where *X* suspends judgment between noncountably many rival hypotheses must be precluded.

It seems to me gratuitously dogmatic to follow the second route, and, in any event, doing so flouts our practice. One of the achievements of Kolmogorov and his predecessors was to enable us to deal with such cases with the aid of measure theory—albeit while retaining some restrictions of doubtful legitimacy, such as countable additivity. In any case, that approach does allow credal regularity to be violated and it would be desirable to follow this view.

Let there be *n* hypotheses exclusive and exhaustive relative to *K*. Inductive logic should not mandate assigning equal *Q*-values of $1/n$ to all *n* hypotheses; but it should not forbid a credal state which does so.

Case 2: The same is true, as *n* goes to infinity and we face a countably infinite number of alternatives. We should not obligate *X* to assign each alternative equal *Q*-value; but we should not forbid it. But we cannot assign positive *Q*-values to all the alternatives which are equal without violating the condition of finite additivity or normalization. Hence, credal coherence mandates that all hypotheses must be assigned 0 *Q*-value.

As a variant of this case, consider a situation where *X* is interested in the value of a continuous variable taking real values between $-\infty$ and $+\infty$. If we wished to regard all point estimates as equally probable in this sort of case, not only would each point estimate be assigned 0 probability, but the point estimate would be assigned 0 density so that any hypothesis to the effect that the true value fell in a given finite interval would also bear 0 probability. Again, a credal state represented by *Q*-distributions of this sort should not be mandated; but neither should it be prohibited.

Case 3: Suppose it is known that a coin is to be tossed until it lands heads for the first time. It is also known that the coin is fair. Thus, *X*'s degree of credence that the coin tossing will stop on the *n*th toss is $(.5)^n$. The hypothesis that the tossing will never stop is, however, a serious possibility (if we can assume some mechanism is operative which secures that the conditions specified above do, indeed, hold). Yet, it receives 0 *Q*-value.

Such a case should not be precluded by principles of inductive logic.

Keep in mind that I am not claiming that in realistic situations we will ever find cases which we must literally describe as being of one of these kinds. Perhaps the coin tossing game cannot be fully implemented. Perhaps no one ever seriously allows it to be possible that the log of the variance of some normally distributed random variable takes any value between

$-\infty$ and $+\infty$. We may never be compelled to face such cases; but we should not rule them out on grounds of inductive logic.

Nonetheless, there is a genuine difficulty. Shimony has a strong argument in favor of credal regularity. REU and *WD* imply it.

Should we follow De Finetti's example and give up *WD*?

I think not. Ramsey,⁶ De Finetti,⁷ and Shimony⁸ all agree that conditional *Q*-values can also be employed in ranking options with respect to expected utility in order to discount what I shall call irrelevant possibilities. And the method of ranking with respect to expected utility so obtained contradicts REU in its unmodified form.

Furthermore, this contradiction emerges only when credal regularity is violated.

Consequently, if we are to refuse to endorse credal regularity, we must avoid trouble on two fronts: the violation of *WD* and the conflict between REU and the alternative method of ranking. This can be done with one fell swoop by abandoning REU.

The critical element in this argument is how conditional *Q*-functions play a role in ranking options with respect to expected utility and, hence, in the determination of *E*-admissibility.

5.6 Irrelevant Possibility and Conditional Probability

Consider the decision problem in classical form shown in table 5.1. The matrix entries are utility values. *W* and *L* are positive.

If the options are credally irrelevant to the states, it seems entirely plausible that the serious possibility that *e* is false (i.e., that $\sim e$ is true) may nonetheless be ignored in ranking the two options with respect to expected utility. This holds for the comparison of *A_i* and *A_{i'}*, even if they are just two of many feasible options in the given decision problem and even if the payoff structure according to the *u*-function is different for the columns under *h & e* and $\sim h \& e$.

In general, in a decision problem where the options are credally irrelevant to the states, the state *h_k* is not a *relevant possibility for the purpose of comparing A_i and A_{i'}* (even though it is a serious possibility) relative to a given *u*-function if and only if $u_{ik} = u_{i'k}$.

In cases such as those described by the matrix representation, where $\sim e$ is an irrelevant possibility, Ramsey, De Finetti,

Table 5.1 Decision Matrix

Options	States of Nature		
	<i>h & e</i>	$\sim h \& e$	$\sim e$
<i>A_i</i>	<i>W</i>	$\sim L$	<i>K</i>
<i>A_{i'}</i>	0	0	<i>K</i>

The matrix entries are utility values. *W* and *L* are positive.

and Shimony recommended ranking *A_i* over *A_{i'}*, with respect to expected utility relative to *Q* and *u* if and only if $Q(h; e)/Q(\sim h; e) > L/W$. But this holds if and only if the *conditional expected utility* of *A_i* on *e* is greater than the conditional expected utility of *A_{i'}* on *e*.

The conditional expected utility of an option *A_i* on *e* in a classical decision problem is given by:

$$E(A_i; e; Q, u) = \sum_{k=1}^n Q(h_k; e \& A_i \text{ is implemented})u_{ik}.$$

Under the condition of credal irrelevance, the options are *Q*-independent of the states. Hence, if *e* is a disjunction of some subset of *U* and *h_k* is a disjunct of *e*,

$$\begin{aligned} Q(h_k; e \& A_i)Q(e; A_i \text{ is implemented}) \\ &= Q(h_k \& e; A_i \text{ is implemented}) \\ &= Q(h_k; A_i \text{ is implemented}) \\ &= Q(h_k). \end{aligned}$$

Moreover, $Q(e; A_i) = Q(e)$. But

$$Q(h_k; e)Q(e) = Q(h_k \& e) = Q(h_k).$$

Hence,

$$Q(h_k; e \& A_i \text{ is implemented}) = Q(h_k; e).$$

Thus, we have

$$E(A_i; e; Q, u) = \sum_{k=1}^n Q(h_k; e)u_{ik}.$$

In the sequel, I shall write $E(A_i; e)$ in lieu of $E(A_i; e; Q, u)$ where no misunderstanding is likely to arise.

In cases where credal irrelevance obtains in a decision problem, and letting *e** be the disjunction of all those elements in *u* which are not irrelevant possibilities for comparing *A_i* with

A_i , then Ramsey, De Finetti, and Shimony appear committed to ranking these two options in this way:

A_i outranks $A_{i'}$ if and only if $E(A_i; e^*) > E(A_{i'}; e^*)$.

At the same time, these same authors are also committed to ranking the same options in accordance with the unconditional expected utility function or in accordance with the conditional expected utility conditional on e^{**} , where e^{**} is a disjunction of elements of U that includes all states which are relevant possibilities for the purpose of comparing A_i and $A_{i'}$ but that may also include some irrelevant possibilities as well.

In effect, they are committed to the following method for ranking feasible options with respect to expected utility relative to Q and u when options are Q -independent of states relative to Q :

Ranking with Respect to Conditional Expected Utility (RCEU): Let e^{**} be a disjunction of elements of U including all elements of U relevant for comparing A_i and $A_{i'}$. A_i outranks $A_{i'}$ if and only if $E(A_i; e^{**}) > E(A_{i'}; e^{**})$.

The disjunction of all elements of U contains all relevant possibilities. Hence, RCEU entails REU.

On the assumption that the ranking obtained with respect to Q and u is a consistent one, the following condition must be satisfied:

Strong Condition of Compatibility of Evaluations of Expected Utility: Let e^* be the disjunction of all and only elements of U relevant for comparing A_i and $A_{i'}$ and e^{**} a disjunction of elements of U entailed by e^* . Then

$E(A_i; e^*) > E(A_{i'}; e^*)$ if and only if $E(A_i; e^{**}) > E(A_{i'}; e^{**})$.

If credal regularity is satisfied, the strong condition of compatibility can be guaranteed to hold. It can also be guaranteed to hold on the assumption of credal coherence provided that we modify condition (1) of section 4.2 so that E is the set of all hypotheses in L bearing positive unconditional Q -values. Let us call this condition (1*).

Suppose that in the matrix for A_i and $A_{i'}$, $Q(h \& e) = Q(\sim h \& e) = Q(e) = 0$. $E(A_i) = E(A_{i'}) = K$. On the other hand, it could very well be the case that $E(A_i; e) > E(A_{i'}; e)$.

But RCEU mandates that the pair of options be ranked in accordance with both measures of expected utility and, hence, expects them both to rank in the same manner as the strong condition of compatibility requires. This is not the case in our example. Hence, either we abandon RCEU, impose credal regularity, or refuse to define $Q(h; e)$ and $Q(\sim h; e)$ in this case and, hence, avoid defining $E(A_i; e)$ and $E(A_{i'}; e)$.

I have already indicated that imposing credal regularity is objectionable in my view. I follow De Finetti also in refusing to restrict the set E in condition (1) of section 4.2 to hypotheses bearing positive Q -value.⁹ The reasons for this are substantially the same as those which favor refusing to impose credal regularity.

Consider a case 1 situation from the preceding section. In particular, let the situation be one where the rival hypotheses are point estimates of the unknown chance of obtaining heads on a toss. Let h_p assert that the chance of r heads in n tosses is $\binom{n}{r} p^r (1-p)^{n-r}$. If X knows that a coin is to be tossed n times but does not know whether h_p is true, and if e_r is the hypothesis that on that particular sequence of tosses the coin will land heads r times, $Q(e_r; h_p)$ is often taken to be equal to the chance of r heads in n tosses when h_p is true—i.e., $\binom{n}{r} p^r (1-p)^{n-r}$. This is due to the principle of direct inference, to be discussed in chapter 12. The point to be emphasized here is that $Q(e_r; h_p)$ is defined even though $Q(h_p) = 0$. Moreover, it is customary practice in statistical applications to proceed in this manner even though the Kolmogorovian axioms on probability prohibit the definition.

Case 2 also provides examples. If we are given a countable infinity of exclusive and exhaustive alternatives and are to take them all as bearing equal Q -value (which must, therefore, be 0), we may wish to assign $Q(h_k; e)$ the value $1/n$ where e is a disjunction of n elements from these alternatives including h_k .

In my opinion, once we allow violations of credal regularity, we should not balk at allowing conditional probabilities to be defined on hypotheses with 0 Q -values. It appears that condition (1) imposes the right requirements for membership in E . $Q(h; e)$ is defined in case e is consistent with the corpus—i.e., its truth is a serious possibility according to the corpus.

Thus, to avoid the contradiction, either credal regularity

should be imposed or the strong compatibility requirement abandoned. I favor the latter alternative.

Given the conditions on Q -functions introduced in section 4.4 and required by credal regularity, there is a weaker compatibility requirement which may be imposed:

Compatibility of Evaluations of Expected Utility: If e^* is a disjunction of states in U including all those which are relevant possibilities for comparing A_i with $A_{i'}$ relative to u , if e^* entails e^{**} , and if $E(A_i; e^{**}) > E(A_{i'}; e^{**})$, then $E(A_i; e^*) > E(A_{i'}; e^*)$.

Abandoning RCEU and restricting the ranking procedure to REU avoids entailing the strong compatibility requirement while conforming to the compatibility requirement. Indeed, the latter requirement then plays no role in ranking with respect to expected utility; for conditional expectations are never used.

This, in my opinion, is a serious objection to this approach; for conditional Q -values lose the significance in guiding appraisals of options in decision making which De Finetti, Ramsey, and Shimony all agreed they should have. Once that function is lost, it becomes unclear why conditional Q -functions should obey requirements (1)–(4) of section 4.2 or why the multiplication theorem should be satisfied.

The compatibility condition acquires teeth, however, if we do wish to follow Ramsey, De Finetti, and Shimony in allowing conditional expectations a role in ranking options with respect to expected utility.

To be sure, RCEU must still be modified. But there are alternatives to reverting to REU. One can devise some method of ranking (several can be envisaged) which gives precedence to rankings in accordance with conditional expected utility, where the condition is a disjunction of all and only relevant possibilities over rankings with respect to unconditional expected utility or conditions entailed by the disjunction of all and only relevant possibilities.

A specimen of such a method of ranking with respect to expected utility is the following:

Qualified Ranking with Respect to Expected Utility (QRCEU):¹⁰ Let e^* be the disjunction of all states of nature

in U relevant for comparing A_i and $A_{i'}$. A_i ranks over $A_{i'}$ if and only if there is a disjunction e^{**} of elements of U entailed by e^* such that $E(A_i; e^{**}) > E(A_{i'}; e^{**})$.

QRCEU induces a quasi ordering of feasible options with respect to expected utility relative to Q and u . The quasi ordering becomes an ordering when the Q -function satisfies regularity. In that case, the ordering coincides with that induced by RCEU and REU.†

One can still obtain from such a quasi ordering a set of options optimal relative to Q and u and, hence, proceed to define E -admissibility along the lines indicated in chapter 4.

There is a modification of QRCEU, MQRCEU, that also conforms to the compatibility condition but yields a connected ordering for all options.¹¹ It is not necessary, however, to explore this ordering here.

What is crucial to notice is that if one uses QRCEU, MQRCEU, or some other variant that ranks options as REU and RCEU do when credal regularity obtains, not only do we avoid the contradiction bred by violating credal regularity when the strong compatibility requirement is adopted but *credal regularity may be violated while obeying WD*.

Consider the decision-problem matrix given in table 5.1. Modify the payoffs by allowing L to be negative and, hence, $-L$ to be positive. In that case A_i weakly dominates $A_{i'}$. If $Q(e) = 0$, $E(A_i) = E(A_{i'}) = K$. However, given conformity to conditions (1)–(4) on Q -functions of section 4.2, then $E(A_i; e) > E(A_{i'}; e)$. There is a conflict in the ranking according to unconditional and conditional expected utility. The ranking according to unconditional expected utility violates *WD*. The ranking according to conditional expected utility does not. QRCEU and MQRCEU both give precedence to the conditional ranking.

Thus, the argument which justifies imposition of credal regularity in order to secure conformity to *WD* is less than compelling not because Shimony or others who have used it made mistakes in their proofs but because they smuggled in debatable assumptions—in particular, assumptions to the effect that

† T. Seidenfeld has pointed out to me that QRCEU fails to induce an ordering. H. Stein showed that it would yield a quasi ordering.

rankings with respect to expected utility should agree with REU.

The considerations adduced here suggest that there are good reasons to reject this kind of assumption. We need not decide between allowing violations of credal regularity and requiring conformity with *WD*; and, in the light of the considerations adduced here, it seems sensible to do both.

5.7 Countable Additivity

Case 2 has been used previously to illustrate situations where credal regularity is violated and where conditional probability on conditions with 0 *Q*-values are defined.

It also illustrates situations where countable additivity is violated. Countable additivity requires that if *g* asserts that one of a countable infinity of exclusive and exhaustive hypotheses is true, then $Q(g; e) = \sum_{i=1}^{\infty} Q(h_i; e)$ where $h_i \in U =$ the set of exclusive and exhaustive hypotheses (relative to *K*).

Given countable additivity,

$$Q(g \& f; e) = \sum_{i=1}^{\infty} Q(h_i \& f; e) = \sum_{i=1}^{\infty} Q(f; h_i \& e)Q(h_i; e).$$

Hence, if $K, e \vdash g$, then $K, e \vdash g \& f \equiv f$. Therefore,

$$Q(f; e) = \sum_{i=1}^{\infty} Q(f; h_i \& e)Q(h_i; e).$$

This last result captures the force of what De Finetti calls *conglomerability*.¹² If we do not impose countable additivity as a requirement on *Q*-functions, conglomerability is not required either.

Once more, I follow De Finetti in refusing to adopt a requirement of countable additivity because it prohibits as a matter of inductive logic the adoption of credal states which ought not to be prohibited in this categorical manner.

Thus, requirements to be imposed on *Q*-functions by the principle of credal coherence ought, in my opinion, to be restricted to those introduced in section 4.2.

5.8 Credal Consistency

Credal consistency requires that a credal state relative to *K* should be nonempty if and only if *K* is inconsistent. The rationale of this requirement is that if, relative to consistent *K*, *B* is empty, any decision problem *X* faces when *K* is his

corpus can have no *E*-admissible options and, by the criteria for admissibility to be constructed, can have no admissible options. The implication would be that even if *X* is constrained to choose one of the options (although he is free to choose which one he will implement) there is no option which *X*, as a rational agent, is permitted to choose.

Any theory of rational choice which leads to such a result should be revised. Credal consistency, in the light of the other requirements, prevents such an untoward consequence.

Of course, the rationale for valuational consistency is precisely the same; for an empty *G* will also lead to the absence of *E*-admissible options.

Of the conditions on credal rationality, only the requirement of credal convexity remains to be considered. That topic will be postponed until chapter 9, when sufficient background will be available for consideration of various aspects of the matter.

5.9 On Infinitely Many Alternatives

Suppose *X* is concerned to find out which of *n* exclusive and exhaustive alternatives $h_1, h_2, \dots, h_n \in U$ are true relative to corpus *K*. For the present we are concerned with properties of probability measures and, hence, it will be convenient to proceed as though *X*'s credal state contained exactly one *Q*-function. This may happen but need not and should not on all occasions. But under the circumstances, it is convenient to focus on those occasions where credal uniqueness obtains.

Whether the *Q*-function satisfies credal regularity or not, given the unconditional *Q*-values for all elements of *U*, the unconditional *Q*-values for all hypotheses equivalent given *K* to disjunctions of finitely many elements of *U* are all determined.

Furthermore, if credal regularity obtains, the *Q*-function (defined over a Boolean algebra) conditional on the truth of the disjunction of finitely many elements of *U* can always be derived from the unconditional *Q*-function for the same algebra.

Suppose, however, that credal regularity breaks down. Then unconditional *Q*-functions cannot determine conditional ones in all cases. This is so even when *U* is finite.

Thus, if *U* consists of three hypotheses h_1, h_2, h_3 , where h_1 and h_2 bear 0 *Q*-values and h_3 bears the *Q*-value 1, we cannot

automatically determine the value of $Q(h_i; h_1 \vee h_2)$ from the multiplication theorem.

It may be thought that when U is finite there is no pressure to consider irregular Q -functions. But that is by no means always so.

Suppose X wishes to test the hypothesis that the leaf color of some plant is inherited in accordance with a simple Mendelian model. X knows that he is crossing plants which are purebreds with respect to traits A and a , and he also knows that A dominates a . (Past breeding experiments might have established this.) If the simple Mendelian model is correct, the chance of an offspring exhibiting the dominant trait should be exactly $\frac{3}{4}$.

But X initially might suspend judgment as to what the true chance is. As far as he is concerned, the true chance might be any real value between 0 and 1. Moreover, X might assign 0 credence to every point estimate of the value of that chance.

Consequently, even though in the context of his investigation X might be concerned only with whether the chance is $\frac{3}{4}$ or not and, hence, with a finite U , credal regularity will be violated.

Another example is to be found in cases where X wishes to estimate the unknown chance in a binomial process (as in the above example) but where U consists of the three hypotheses: h_1 ($p = 0$), h_2 ($p = 1$), and h_3 ($0 < p < 1$). Here the first two alternatives bear 0 Q -value, and the third a Q -value of 1. At any rate, a credal state of this sort ought not to be prohibited as irrational. And in this case, the value of $Q(h_i; h_1 \vee h_2)$ is undetermined.

This does not mean that this conditional Q -value is undefined, but only that it cannot be derived from Q -values obtained unconditionally for the elements of U and Boolean combinations thereof.

In my opinion, this result is inescapable and quite acceptable.

Our examples, however, do point to a moral which ought to be emphasized. In most cases where there is pressure to violate credal regularity when U is finite, it is because there is some refinement of U into a system of exclusive and exhaustive alternatives which is infinite. In the examples we

have been considering, indeed, the new U is uncountably infinite.

Here, however, we face several technical problems. Consider first those cases where U is countably infinite. For them, our regimented language L may be able to express all the elements of U by sentences $h_1, h_2, \dots, h_i, \dots$. However, it will be desirable to assign Q -values to all hypotheses which assert that the true h_i belongs to a subset of U . There are noncountably many such hypotheses and at most countably many of these are expressible in L .

There are several ways to handle this matter. I shall proceed by appending indices to the elements of U , and then introduce a probability measure defined over all elements of the power set of all the indices. Countably many of these sets of indices that are members of the power set will be correlated with sentences in L asserting that the true h_i has an index in the power set; and these sentences will obtain Q -values equal to the probability measure assigned the corresponding set.¹³

The power set of the set of indices is not merely a Boolean algebra but a σ -algebra and the underlying probability measure defined over the power set (which shall, when no confusion threatens, be called a Q -function just like the measure on the sentences in L) is, therefore, a measure on a σ -algebra.

The unconditional probability measure on the σ -algebra may obey credal regularity or it may violate it. It may be countably additive or fail to be so.

When countable additivity obtains, the situation appears analogous to that which obtains in the finite case. Whether credal regularity holds or not, the unconditional Q -function for elements of U uniquely determines the unconditional Q -function over the entire σ -algebra. If credal regularity is satisfied, the unconditional Q -function over the σ -algebra uniquely determines the conditional Q -values for all elements of the algebra conditional on any member of the algebra consistent with K . When credal regularity is violated, this no longer obtains.

When countable additivity fails, the situation changes. Before considering such cases, however, let us first consider situations where U is noncountably infinite.

5.10 The Continuous Case

Suppose X wishes to estimate the true value of θ where K asserts that θ is some real number in the interval from $-\infty$ to $+\infty$ or asserts that it falls in some subinterval thereof. (The argument will be the same if K asserts that θ is an n -tuple of real numbers located in an n -dimensional space or in some specified region thereof.) U is then the set of point estimates and is noncountably infinite. For the purpose of this essay, attention will be focused on situations where noncountably infinite U is of this kind or can readily be handled once this kind of case is understood.

In such cases, not even the elements of U can be expressed in L . We have more reason than before to begin with an underlying probability measure defined for points and sets of points in the given n -dimensional space and then use that measure to induce Q -values for hypotheses in L which assert that the true value of θ belongs to such and such a set.

We are supposing that length of a (closed, open, half open) interval in the one-dimensional case (or the volume of an n -dimensional interval in the n -dimensional case) is defined and that this measure is extended to the σ -algebra of either Borel sets or Lebesgue measurable sets of elements of U .

A probability measure can be induced over the elements of this σ -algebra and then the question might arise as to whether it can be extended to apply to all sets in the power set of U .

If the extended measure is to obey countable additivity, the extension cannot be implemented. Hence, either one thinks of restricting the domain of definition to the domain of Lebesgue measurable sets or gives up countable additivity. De Finetti argues for the latter alternative and I sympathize.¹⁴

However, one could, nonetheless, retain a qualified endorsement of countable additivity of the following kind. One could insist that countable additivity apply to those sets that are Lebesgue measurable, although it must be violated in the extension. The considerations arguing against countable additivity alluded to previously suggest that this requirement remains excessively restrictive. Nonetheless, many probability measures employed in applications do satisfy the requirement and it is worth considering some of the philosophically interesting features. Consequently, when I say that countable additivity is satisfied in the continuous case, I mean that countable additivity is satisfied by the measure on the domain of

Lebesgue measurable sets but not on the domain of elements of the power set. When, in the next section, I say that countable additivity is allowed to be violated, I mean that it is violated even within the domain of the σ -algebra of Lebesgue measurable or even Borel sets.

Under the circumstances specified, at most countably many points in U may bear positive unconditional probability. Noncountably many such points must bear 0 unconditional Q -values. Credal regularity is perforce violated.

Furthermore, even if countable additivity is satisfied, the Q -values for all Lebesgue measurable sets are not uniquely determined by the Q -values for the points in U . In general, a specification of the unconditional Q -values for all n -dimensional volumes (including degenerate volumes of intervals of lower dimension) must be given.

In the *continuous cases* to be considered here, Q -values can be determined by starting with a *cumulative distribution function* (cdf) $F(\theta) = F(x_1, x_2, \dots, x_n)$ which specifies the Q -value for the hypothesis that the i th component of θ has a value less than or equal to x_i for each i and where the cdf $F(\theta)$ is continuous and totally differentiable over the domain of points in U . *Discrete cases* where there are at most countably many points of discontinuity can also be treated, but attention will be focused on the continuous cases here.

An alternative mode of representation is obtained by using the *density function* $f(x_1, x_2, \dots, x_n) = \partial^n F / \partial x_1 \partial x_2, \dots, \partial x_n$.

These are all alternative methods of characterizing the (unconditional) *joint distribution* over the points in the n -dimensional space and each can be used to uniquely determine the Q -value associated with each measurable set. It is important to remember, however, that all of this presupposes countable additivity over the domain of Lebesgue measurable sets.

Because credal regularity is violated, unconditional Q -values will often fail to determine conditional Q -values. On the other hand, if countable additivity is satisfied, there is an extension of the multiplication theorem that may be used to partially—but only partially—remedy the situation.

For notational simplicity, let us suppose the values of θ to be points in a two-dimensional space. The observations to be made here can be extended to higher dimensions if desired. Let $f(x, y) = f(\theta)$ be the *joint density function*, which equals

$\partial^2 F(x, y)/\partial x \partial y$. If countable additivity holds, we can define a *marginal* cumulative distribution function $G(x)$ for x specifying the Q -value that the value of that real variable is less than or equal to x ; for it should be equal to the Q -value for that hypothesis conditional on the value of y falling anywhere in the range allowed by the corpus K and this, in turn, is equal to $F(x, \bar{y}) - F(x, \underline{y})$ where \bar{y} is the maximum value permitted to y and \underline{y} the minimum. (These could be $+\infty$ and $-\infty$ respectively.) Then

$$g(x) = dG/dx = \partial F(x, \bar{y})/\partial x - \partial F(x, \underline{y})/\partial x = \int_{\underline{y}}^{\bar{y}} f(x, y) dy.$$

The marginal distribution for y can also be characterized by a marginal cdf $H(y)$ and a marginal density $h(y)$.

For some specific value x^* , let $g(x^*) > 0$. Then $g(x^*) dx$ approximates the Q -value for the hypothesis that the true value of x falls between x^* and $x^* + dx$; this Q -value is positive. $f(x^*, y) dx dy$ specifies that the true value of x falls between x^* and $x^* + dx$ and that the true value of y falls between y and $y + dy$. In virtue of the multiplication theorem, the conditional probability that the value of y falls in its specified interval conditional on the assumption that the value of x falls in its given interval is approximated by

$$f(x^*, y) dx dy / g(x^*) dx = f(x^*, y) dy / g(x^*).$$

Now this ratio performs the function of a conditional density function for y , and can be used to compute Q -values for hypotheses that y falls in some interval conditional on the true value of x being in the interval from x^* to $x^* + dx$. At least, this is so to a good degree of approximation.

What about the distribution for y conditional on the true value of x being x^* ? Here we have a case where the condition bears 0 Q -value and the multiplication theorem mandates nothing. However, because the function $h(y; x^*) = f(x^*, y) / g(x^*)$ is the limiting case of conditional densities where the condition is that x falls between x^* and $x^* + dx$ for decreasing dx , the multiplication theorem is extended to govern such densities so that $f(x, y) = h(y; x) g(x)$.

There is nothing in our original conditions which mandates adopting a credal state where this extended version of the

multiplication theorem obtains. However, where countable additivity obtains in the continuous case, in normal applications it is taken for granted that it does hold and I shall consider only such cases. In this way, the capacity to compute conditional Q -functions from unconditional ones is somewhat extended.

Care must be taken, however, not to use this extension beyond its permitted scope of applicability. Consider, for example, the hypothesis that the true value of x is either x^* or x^{**} . That hypothesis has unconditional credence of 0. The conditional densities $h(y; x^*)$ and $h(y; x^{**})$ are both determined by the argument just given. What about $h(y; x^* \vee x^{**})$? One might think that the analogy with probability computations could be extended further and this density equated with $h(y; x^*) Q(x^*; x^* \vee x^{**}) + h(y; x^{**}) Q(x^{**}; x^* \vee x^{**})$. This much is correct. But what about $Q(x^*; x^* \vee x^{**})$? Nothing in what has been said warrants equating this with $g(x^*) / (g(x^*) + g(x^{**}))$.

To be sure, nothing precludes assigning the Q -values in this way. However, by suitable transformation of x (e.g., to $\log x = z$) we obtain a density function $k(z) dz = g(x) dx$ so that $k(z) = g(x) x$. $Q(z^*; z^* \vee z^{**})$ must equal $Q(x^*; x^* \vee x^{**})$. But $g(x^*) / (g(x^*) + g(x^{**})) \neq k(z^*) / (k(z^*) + k(z^{**}))$.

Since there is no useful way of identifying which variable one should use to compute conditional Q -values from unconditional densities, there is no reason to restrict attention to conditional Q -values constructed one way or another.

Nonetheless, the foregoing method for extending the multiplication theorem with the aid of joint and marginal densities does have useful applications. It appears in standard applications of Bayes' theorem. I shall take for granted that it is operative unless explicit indication to the contrary is given.

5.11 Violating Countable Additivity

Occasions where countable additivity is violated can arise both when U is countably infinite and when it consists of noncountably many hypotheses representable by points in a region of an n -dimensional space. As noted before, stock examples are, in the countable case, assigning equal Q -value to all elements of a countably infinite U and, in the continuous case, assigning equal Q -value to all line segments of equal length of the real line from $-\infty$ to $+\infty$. If the assignments in

question are all positive, *finite* additivity entails that condition (4) of section 4.2 cannot be satisfied.

De Finetti rightly concludes from this that the Q -values in question should be set at 0 if credal coherence is to be satisfied as everyone with any sympathy for Bayesian ideas requires. But, in that case, countable additivity is violated.

The customary response is to endorse countable additivity and to prohibit credal states of the sort just cited. Curiously enough those who are sympathetic to De Finetti's view that such prohibition is unwarranted have not been prepared always to bite the bullet and abandon countable additivity.

H. Jeffreys, for example, preferred assigning positive Q -values to elements of U when they are countably infinite and also positive density to points on the real line in the case of the uniform distribution, even though for both these assignments the total probability assigned to all elements of U (to the real line) then becomes infinite.¹⁵ Jeffreys excused his procedure by suggesting that it constituted a mere change in the scale on which probabilities are measured.

It is, indeed, true that condition (4) of section 4.2 can be altered by substituting any other positive finite number for 1 without changing the account of credal rationality in any important way—provided that the same positive finite number is used as the maximum for probability throughout. In this sense, the choice of 1 is a convention as Jeffreys rightly insists.

But Jeffreys' procedure cannot be excused by this observation. In the first place, if the maximum probability is infinite, no change of scale (i.e., no method of multiplying the probability measure by a positive finite constant) will diminish it to a finite value. Furthermore, as I. Hacking has pointed out, the same maximum should be used throughout one's calculations.¹⁶ Jeffreys uses 1 as the maximum except when compelled to do otherwise in order to represent uniform distributions.

There is no escaping the conclusion drawn by De Finetti that if one wishes to allow credal states of the sort of interest to Jeffreys (and De Finetti agrees with Jeffreys in wishing to do so) the elements of U must be assigned 0 Q -value or 0 density.

In spite of this, Jeffreys' instincts are, in my opinion, quite sound.

Given an assignment of 0 as the Q -value for each element of U in the countable case or an assignment of 0 as the Q -value for every finite interval in the continuous case, there is no determination of the Q -values for the other elements of the σ -algebra of members of the power set when U is countably infinite or of Lebesgue measurable sets in the continuous case. Furthermore, assigning 0 Q -value fails to make discriminations which there is some warrant for seeking to represent.

Consider a countably infinite U and contrast a case (i) where each h_i in U is intended to have equal Q -value with every other and where the conditional Q -value for h_i conditional on one of n elements of U (including h_i) being true is $1/n$ with the case (ii) where U^* is obtained from U by refining the partition through identifying g_{11} with h_1 , construing h_2 as the disjunction of g_{21} and g_{22} , and, more generally, construing h_i as the disjunction of i exclusive alternatives $g_{i1}, g_{i2}, \dots, g_{ii}$ where $Q(g_{ij}; h_i) = 1/i$. In this latter case, $Q(g_{11}) = Q(g_{21}) = 0$. Yet, there is a sense in which the evaluations of these alternatives with respect to credal probability do not rank them equally.

One way to bring this out is to define $Q(g_{11}; g_{11} \vee g_{21}) = 2/3$. Observe, however, that this assignment is not mandated by those previously made. And, in any case, insofar as feasible it would be desirable to be able to determine such conditional Q -values from unconditional appraisals.

Jeffreys was seeking ways and means to do this and, in so doing, for extending the applicability of Bayes' theorem to situations where "prior" Q -distributions are improper in the sense we have been considering.

His mistake, in my opinion, was not in using a measure function which obeys countable additivity but "blows up" because it assigns positive measure to infinitely many alternatives covering the entire space of alternatives. Rather, he erred in taking that measure to be the probability measure.

We have already seen that there are other ways to represent credal states than by means of probability measures. Q -functions can be determined, for example, by densities or cumulative distribution functions. It is a mistake, however, to think of a density measure as a probability measure. So too it is a mistake to think of Jeffreys' "improper" measures as probability measures. They do, however, have a use in the systematic representation of Q -functions.

Given a countably infinite U , consider the power set generated by U . An (unconditional) measure on U is a σ -finite measure defined for all members of the power set if and only if it is (i) nonnegative, (ii) countably and, hence, finitely additive, and (iii) there is some countably infinite partition of the set U such that each element of the partition is assigned finite measure.¹⁷ The same definition applies in the continuous case except that the σ -algebra over which the measure is defined is the class of Lebesgue measurable or Borel sets. Otherwise we can obtain density functions as before.

Thus, reverting to our example of countably infinite alternatives, in the case of U^* , we assign g_{11} the σ -finite value $m_K(g_{11}) = k$, $m_K(g_{21}) = m_K(g_{22}) = k/2, \dots, m_K(g_{ij}) = k/i, \dots$. For each h_i in U , we have, on the other hand, $m_K(h_i) = k$.

It does not matter here what positive finite value is attributed to k . Any system which assigns αk instead of k ($\alpha > 0$) will lead to equivalent results as far as we are concerned here.

Given such a σ -finite measure, unconditional Q -values can be determined for many (though not all) elements of the σ -algebra. The rules for the procedure are as follows:

- (a) If $m_K(h)$ and $m_K(\sim h)$ are both finite, then let $Q(h) = m_K(h)/(m_K(h) + m_K(\sim h))$.
- (b) If $m_K(h)$ is finite and $m_K(\sim h)$ infinite, then let $Q(h) = 0$.
- (c) If $m_K(h)$ is infinite and $m_K(\sim h)$ is finite, then let $Q(h) = 1$.

Thus, in our example, $Q(g_{11}) = Q(g_{21}) = Q(h_2) = 0$ even though the m -values for these alternatives are not all the same.

Thus far, we have failed to consider cases where both h and $\sim h$ are assigned infinite m_K -values. Consider, for example, the case where h is associated with the set of alternatives h_i with even subscripts and $\sim h$ with those bearing odd subscripts.

The σ -finite measure we have used will prove insufficient to handle this case. In such cases, we repartition the original partition U or U^* into finitely or countably many alternatives each of which bears an infinite m_K -value and define a new σ -finite measure over these which can then, in turn, be used to determine unconditional Q -values for more hypotheses in the σ -algebra.

This procedure can be reiterated as necessary, introducing as many σ -finite measures as desired.

Thus far only unconditional m_K -values and Q -values have been considered. Let e be consistent with K and consider the conditional measure $m_{K,e}$ associated with m_K . There are three cases to consider:

$m_K(e)$ is finite and positive. $m_{K,e}(g) = m_K(g \& e)/m_K(e) = Q(g; e)$.

$m_K(e) = 0$. $m_{K,e}(g)$ is given independently. It is a σ -finite measure for the partition U' obtained by deleting elements of U inconsistent with K and e .

$m_K(e) = \infty$. $m_{K,e}(g) = km_K(g \& e)$ for some finite constant $k > 0$.

Rules (a), (b), and (c) are to be used to obtain $Q(g; e)$ from $m_{K,e}(g)$.

The most important case is the first in situations where e is equivalent given K to an infinity of alternatives in U .

Thus, let U consist of alternatives of the form e_i & h_j for i and j ranging over all the integers. Let $m_K(h_j)$ be a positive finite constant the same for all h_j 's. Let $m_K(e_i) > 0$. $Q(e_i; h_j)$ and $Q(h_j; e_i)$ are both defined even though $Q(h_j) = 0$ and, perhaps, $Q(e_i) = 0$ as well. Indeed,

$$Q(h_j; e_i) = \frac{Q(e_i; h_j)m_K(h_j)}{\sum_{j=1}^{\infty} Q(e_i; h_j)m_K(h_j)} = \frac{Q(e_i; h_j)}{\sum_{j=1}^{\infty} Q(e_i; h_j)}$$

Thus our rules allow in such cases the use of Jeffreys' extension of Bayes' theorem to derive posteriors from priors. Instead of using a prior probability distribution over the h_j 's, however, the uniform σ -finite measure was used. Had the prior been improper in some other way a different σ -finite measure could have been used.

These techniques can be readily adapted to apply to continuous cases.

This scheme is, in its essentials, derived from A. Renyi.¹⁸ Renyi, however, does not define unconditional Q -values except in cases where both the "event" and its complement bear finite measure according to the σ -finite measure used. Thus, when U is countably infinite and all elements are to be treated alike, this is revealed in the m_K -measure; but Renyi does not introduce a probability measure assigning values to the elements of U . I follow De Finetti in being prepared

to assign Q -values to all elements of the σ -algebra. By failing to do so, Renyi disguises the departure from countable additivity which De Finetti makes explicit.

Yet, the class of cases I am considering here is not quite as extensive as that which De Finetti is prepared to consider. To illustrate, consider the following example.¹⁹

Let U consist of countably many exclusive and exhaustive alternatives such that each h_i is equivalent (given K) to $g_i \vee f_i$, where K entails that $g_i \& f_i$ is false. For each i , $Q(g_i) = Q(f_i) > 0$. Then set $\sum Q(g_i) = \sum Q(f_i) = 1/3$. Let g assert that at least one of the g_i 's is true and f assert that at least one of the f_i 's is true. Both of these hypotheses must receive Q -values at least as great as $1/3$ but one must receive a Q -value greater than $1/3$. Hence, countable additivity will be violated. Yet all the other requirements on Q -functions, including credal regularity, may be satisfied.

This example of De Finetti's cannot be represented with the aid of σ -finite measures of the sort introduced here. In the spirit of tolerance which has been the basis for my defense of violations of countable additivity in the first place, I see no reason for ruling out such cases as a matter of principle. On the other hand, none of the applications where countable additivity is violated—such examples will appear in later chapters—will be of this sort. At present, De Finetti's example is an interesting mathematical possibility. Its significance for applications remains obscure.

In recent years, M. Stone and other statisticians have introduced a raft of paradoxes which arise in cases where "improper" distributions of the sort utilized by Jeffreys are introduced. It is argued that the "contradictions" which emerge provide decisive reason for prohibiting such improper distributions. As a consequence, important practitioners of this black art such as D. V. Lindley have recanted.²⁰

An example of such a paradox will be discussed in section 12.6. In that example and others like it, contradictions result from the use of countable additivity. Since countable additivity is already violated by the introduction of the improper prior (understood as a σ -finite representation of a finitely but not countably additive prior) and since the derivation of a paradox employs countable additivity, it is scarcely surprising that contradiction emerges. The source of the difficulty, however,

is not in the use of improper distributions but in the inconsistent assumption and rejection of countable additivity. De Finetti argued against this practice many years ago. It is high time we took his admonitions to heart.²¹

Many of the paradoxes discussed by Stone are not to be analysed in the manner of section 12.6. T. Seidenfeld argues that these so-called "marginalization" paradoxes concern reasoning about conditional probabilities where the conditions bear 0 probability as well as countable additivity. He intends to discuss this matter elsewhere. In any case, the marginalization paradoxes pose no obstacle to the use of improper priors—i.e., priors violating countable additivity.

6.1 The Rule for Ties

Sections 2.3–2.8 contain a summary of the account of inferential expansion proposed in *Gambling with Truth*¹ and extended in “Information and Inference”² and “Acceptance Revisited.”³ I assume in those sections, as in my previous publications, that the criteria for choosing expansion strategies should be derivable from general principles of rational choice together with a specification of the goals and options involved in inferential expansion.

These accounts presuppose that strict Bayesian conditions on credal rationality are satisfied and that the cognitive goals are representable by an “epistemic utility function” unique up to a linear transformation. This utility function is determined (up to a linear transformation) once the M -function and the value of the caution parameter q are specified.

I have always accepted these strict Bayesian requirements with some embarrassment;⁴ but, in this section, I shall temporarily continue to require them.

If the principle of expected utility, which specifies that E -admissibility is both necessary and sufficient for admissibility, is adopted, it follows that all and only those expansion strategies which are optimal (and hence bear maximum expected epistemic utility) are admissible. However, according to the theory developed in chapter 2 and in my previous work, when two or more expansion strategies are E -admissible, it remains the case that exactly one is admissible.

The cognitive options or expansion strategies may be partially ordered with respect to strength by counting one option stronger than another if it rejects a proper subset of elements of the ultimate partition U containing all elements of U rejected by the second option (or if the disjunction of elements of U “accepted as strongest via induction from K ” according to the first option entails (given K) the disjunction of elements of U accepted as strongest according to the second expansion strategy but not vice versa).

Moreover, I have shown that, given strict Bayesian conditions on credal and valuational rationality, the epistemic utility functions I propose guarantee the existence of exactly one E -admissible expansion strategy such that no other E -admissible expansion strategy is weaker than it is.⁵

Taking this into account I proposed a “rule for ties,” according to which that particular weakest E -admissible cognitive option should be counted uniquely admissible among the cognitive options.⁶

The rule for ties violates the principle of expected utility by denying that E -admissibility is sufficient for admissibility. This ought not, however, count as an objection to the rule for ties.

A satisfactory account of admissibility ought to guarantee that if there is exactly one E -admissible option it should be uniquely admissible. To this extent considerations of expected utility ought to dominate all other considerations in assessing admissibility.

It is tempting to extend this requirement to situations where two or more options are E -admissible, and to insist that in those situations as well only E -admissible options should be admissible. As a rule, I think this stipulation is sound; but as shall emerge in the subsequent discussion, I also think a case is available for allowing some exceptions to the rule.

The rule for ties, in the strict Bayesian case, does prescribe choosing an E -admissible expansion strategy in the context of inferential expansion. But it proscribes alternative expansion strategies even though they are E -admissible. Thus it satisfies the requirements just cited. And, in addition, it satisfies a demand deeply ingrained in the epistemological outlook I am attempting to articulate.

In cognitive decision making in general and in inferential expansion in particular, the options can be partially ordered with respect to strength. Indeed, in inferential expansion, the partial ordering has the properties of a lattice. Consequently, given any set of cognitive options, there is exactly one strongest option as weak as or weaker than all options in the set. In choosing that option, X can be viewed as expressing a sort of suspension of judgment between options in the initial set. Thus, if the initial set consists of the options of accepting g as strongest via induction and of accepting g' as strongest, if X

adopts the option of accepting $g \vee g'$ as strongest, he is not only suspending judgment as to the truth values of g and g' , but also expressing a sort of suspense concerning the merits of g and g' by leaving open the opportunity for subsequent expansion to render a verdict between those two strategies.

The aforementioned deeply ingrained demand is manifested, in my view of inferential expansion, in the prescription to suspend judgment between all E -admissible expansion strategies. In cases where two or more cognitive options are E -admissible, I contend that it would be arbitrary in an objectionable sense to choose one over the other except in a way which leaves open the opportunity for subsequent expansions to settle the matter as a result of further inquiry.

I wish to extend this attitude to other contexts of decision making as well. In inferential expansion, it seems clear enough to say that choosing to accept $g \vee g'$ as strongest via induction leaves open the acceptance of g as strongest (or of g' as strongest) and expresses suspension of judgment between these two cognitive options. But in other decision problems, it is not so easy to characterize the sense in which choosing one option expresses suspension of judgment between others.

In my opinion, the problem reduces to identifying some quasi ordering of the options feasible in the decision problem, which identification is interpretable as ordering the options with respect to strength. Given such an ordering, the attitude I wish to emphasize can be brought to bear in other contexts besides inferential expansion.

I shall return to the issue of extending the idea of ordering options with respect to strength to other contexts subsequently. For present purposes, it suffices to insist that considerations of expected utility ought not to so outweigh other considerations as to lead us to allow all E -admissible options to be admissible and so prevent us from favoring suspension of judgment between E -admissible options in those contexts where it makes sense to do so and is feasible to do so.

The rule for ties is not the only principle that may be invoked with some cogency when considerations of expected utility fail to render a verdict. Minimax principles, among others, have been devised to evaluate options when expected utility fails. I shall discuss such principles in the next chapter.

I mention them here only to emphasize that the idea of modifying the principle of maximizing expected utility (understood as counting E -admissibility as necessary and sufficient for admissibility) has ample precedent.

Thus, the rule for ties represents an attitude favoring suspension of judgment over arbitrary choice when, in cognitive decision making, more than one option is E -admissible. I contend that even though using that principle to give expression to that attitude violates the principle of maximizing expected utility, the violation is not without precedent and, in any event, continues to respect the relevance of expected utility in evaluating options with respect to admissibility by insisting that the options between which judgment is to be suspended should be E -admissible ones.

But the rule for ties has been formulated on the assumption that strict Bayesian conditions are satisfied and for cognitive decision problems concerning inferential expansion. In the next section, I shall continue to deal with inferential expansion but shall relax the assumptions embedded in the strict Bayesian case.

6.2 The Rule for Ties Extended

In the strict Bayesian case, X 's credal state for the elements of the ultimate partition U contains exactly one Q -function. His demands for information are represented by a single M -function. The expected epistemic utility of accepting g as strongest is equal to $Q(g) - qM(g)$. This in turn is equal to $\sum[Q(h_j) - qM(h_j)]$ where the sum is taken over all h_j 's in the ultimate partition that are disjuncts in g .

If g bears maximum expected epistemic utility, then (i) no disjunct h_j which is in U and in g is such that $Q(h_j) - qM(h_j)$ is negative; (ii) for every h_j in the ultimate partition such that $Q(h_j) - qM(h_j) > 0$, h_j is a disjunct in g ; and (iii) elements of U for which $Q(h_j) - qM(h_j) = 0$ may or may not be disjuncts in g .

Clearly there must be a weakest cognitive option bearing maximum expected utility—to wit, the g satisfying (i) and (ii) and containing as disjuncts all elements of U cited in (iii).

Hence, in the strict Bayesian case, the injunction to suspend judgment between all E -admissible options leads to the adoption of an E -admissible option. E -admissibility fails to be suf-

ficient for admissibility according to the rule for ties but it remains necessary.

This no longer holds true when we consider violations of strict Bayesian requirements. To illustrate the point numerically, let U contain $h_1, h_2, h_3,$ and $h_4,$ each bearing an equal M -value of $1/4,$ and let $q = .5;$ so that if there were only one Q -distribution over the elements of $U,$ the rule for ties would lead to rejecting all and only those hypotheses whose Q -values were less than $1/8 = .125.$

Let X 's credal state be the convex hull (the set of weighted averages) of the following two distributions: $Q(h_1) = .7, Q(h_2) = .2, Q(h_3) = .01,$ and $Q(h_4) = .09;$ $Q'(h_1) = .7, Q'(h_2) = .01, Q'(h_3) = .2,$ and $Q'(h_4) = .09.$

According to the Q -function, the uniquely E -admissible option is accepting $h_1 \vee h_2$ as strongest. According to the Q' -function, it is accepting $h_1 \vee h_3$ as strongest. If we consider other weighted averages of these two functions, we can find some for which the optimal cognitive option is accepting h_1 as strongest. However, there is no Q -distribution in the credal state according to which accepting $h_1 \vee h_2 \vee h_3$ as strongest is optimal. Hence, this cognitive option is not E -admissible. Yet, it is the strongest cognitive option as weak as or weaker than all the E -admissible cognitive options.

Thus, if we are to extend the rule for ties to apply to inferential expansion when strict Bayesian conditions are not satisfied, we shall have to abandon the requirement that E -admissibility is necessary for admissibility.†

There are two considerations which mitigate the seriousness of this deviation from the principle of maximizing expected utility:

(a) The strongest potential answer as weak as or weaker than all E -admissible potential answers is E -undominated. That is to say, there is no cognitive option which bears higher expected utility than it relative to every permissible Q -function in the credal state and every permissible utility function in the system of goals and values. All E -admissible options are E -undominated; but the converse does not, in general, hold.

(b) I am not suggesting that rational X should ignore con-

† Thus, I modify the position I took in "Indeterminate Probabilities," *J. Phil.*, v. 71 (1974), p. 410.

siderations of E -admissibility. Rather he should avoid whenever feasible rendering a verdict between E -admissible options by remaining in suspense. Thus, E -admissibility continues to be taken into account even though suspension of judgment entails adopting a cognitive option which is not itself E -admissible.

If one is prepared to accept this violation of the injunction to restrict admissible options to E -admissible ones (as I am prepared to do) for the sake of the rule for ties, the criteria for inductive expansion obtained can be characterized as follows:

Reject h_j in U if and only if for every Q -function in $B,$ $Q(h_j) < qM(h_j).$

We may generalize still further. Thus far we have supposed that X 's epistemic utilities are representable with the aid of a definite value of the index of caution q and a single M -function. Let us, for the moment, keep q fixed but allow all M -functions in a given convex set to be used. The criteria for expansion conform to the rule just cited except that h_j is rejected if and only if $Q(h_j) < qM(h_j)$ for every permissible Q -function and every allowed M -function.

Finally, we may allow X to be undecided as to the index of caution $q.$ In that event, he should use the glb of the range of permitted values for q to assess rejection.

The idea behind this procedure is that an element of the ultimate partition U is to be rejected if and only if it is rejected according to all permissible rankings of cognitive options with respect to expected epistemic utility where the procedures for assessing expected utility employed are those described in chapter 2 and proposed by me in earlier publications.⁷

6.3

P-Admissibility

The extended rule for ties presented in section 6.2 applies only to inferential expansion. Inferential expansion, however, is a species of deliberate decision making—so I have maintained. Consequently, the criteria for rational choice regulating inferential expansion ought to be seen as derivable from fundamental criteria for deliberate decision making applicable not only when the goals are cognitive but even in contexts of noncognitive decision making where the goals are moral, political, economic, or otherwise different from the aims of inferential expansion. For this reason, the criteria for rational

choice invoked in the previous section in contexts of inferential expansion ought to be reformulated in a manner allowing recognition that noncognitive as well as cognitive decision making is subject to these criteria.

Following the terminological practice adumbrated in section 4.8, I shall call those options *P-admissible* which survive the tests discussed in section 6.2 when they are generalized so as to cover a broader range of decision problems. *P-admissibility* will be defined in a way which implies that options prescribed by the extended rule for ties in the context of inferential expansion become uniquely *P-admissible*. But the concept of *P-admissibility* will be applicable in other contexts as well.

I am not certain whether the extended concept of *P-admissibility* has much more than a purely formal significance. To repeat, I introduce it here primarily to establish my contention that my account of inferential expansion sees such decision making as a special case of a more general class of decision problems covering both cognitive and noncognitive decision problems. In the next section, I shall consider some simple-minded examples of practical decision making where considerations of *P-admissibility* might turn out to be relevant in a nontrivial way; but the question remains an open one.

I shall suppose that the options feasible for the agent are subject to a quasi ordering with respect to strength. In the special case of a cognitive decision problem, that quasi ordering is a partial ordering with the properties of a lattice. But I am not supposing that these conditions always obtain. Indeed, in most practical decision problems, I suspect they do not. Moreover, I shall postpone until the next section what such a quasi ordering might mean in the context of a practical decision problem. I wish to focus on formal problems for the present.

In section 6.2, it was noted that in inferential expansion with epistemic utility functions meeting the requirements I have imposed, at least one cognitive option exists which is (a) as weak as or weaker than all *E-admissible* options, (b) no weaker than any other option satisfying (a), and (c) *E-undominated*.

The existence of such an option cannot be guaranteed when the quasi ordering with respect to strength fails to meet the requirements of a lattice structure or the utility functions fail

to behave like epistemic utility functions of the sort I have been using. And even if such an option exists, there could be more than one—unlike the situation in inferential expansion as discussed in section 6.2.

Let us call an option meeting the three conditions just satisfied a *WU-option*. In inferential expansion, there must always be one and only one *WU-option*. The extended rule for ties recommends it as uniquely admissible.

It seems entirely plausible to allow that when there are several *WU-options* the considerations of option preservation, suspension of judgment and strength being invoked fail to render a verdict between them. That is to say, all such options should count as *P-admissible*.

Suppose, however, that there are no *WU-options*. In that event, there is no *E-undominated* option which allows for suspension of judgment between all *E-admissible* options. In that case, considerations of option preservation, suspension of judgment, and of strength have failed to render a verdict. The decision maker is left with as indeterminate a solution as he was when he appraised the options with respect to expected utility. In that case, the set of *P-admissible* options should coincide with the set of *E-admissible* options.

These considerations suggest the following definition of *P-admissibility*:

If there are any *WU-options*, all and only *WU-options* are *P-admissible*. If there are no *WU-options*, all and only *E-admissible* options are *P-admissible*.

6.4 On the Strength of Options

In the preceding section I have outlined a way that cognitive options may be partially ordered with respect to strength in a manner relevant to the evaluation of these options in inferential expansion. Should we seek to extend this idea to contexts of noncognitive decision making?

We might attempt to do so in the following manner: If Jacob is convinced that it is feasible for him to marry Rachel, he is certain that choosing to marry Rachel implies his marrying her. Given such convictions, Jacob is entitled to employ a rule or routine for expansion which adds to his corpus those hypotheses whose truth is entailed by the information that he chooses a given option provided that he chooses that option.

Each feasible option in a decision problem is associated with a set of sentences or propositions to be added to the agent's corpus upon choosing that option. This applies to all kinds of decision making and is not restricted to inferential expansion.

Consequently, feasible options might be ordered with respect to the strength of the information added due to such choosing true.

This is not, however, the way in which cognitive options are ordered with respect to strength in efforts at inferential expansion. Consider the option of adding g to K together with all deductive consequences of g and K not already in K . That option is represented as accepting g as strongest via induction from K .

If X chooses that cognitive option, the strongest proposition he chooses true is "X accepts g as strongest via induction from K ." X does not choose g to be true at all.

In cognitive decision making, the options are ranked with respect to the strongest sentences to be added to K when the options are implemented. They are not ranked with respect to the strongest proposition chosen true.

In point of fact, in cognitive decision making, the strongest proposition chosen true in choosing to accept g as strongest does not imply the strongest proposition chosen true in choosing to accept $g \vee g'$ as strongest. The propositions involved are incompatible (given X 's knowledge). Thus, the two propositions are not comparable with respect to strength if the ordering with respect to strength proceeds in this manner. Precisely the same point applies to the comparisons of options in practical decision problems.

Suppose, for example, Laban, with a shotgun pointed at Jacob's head, compels Jacob to marry Rachel or to marry Leah but not both. Jacob has two options. Given his corpus, choosing "Jacob marries Rachel" implies choosing "Jacob does not marry Leah"; and choosing "Jacob marries Leah" entails choosing "Jacob does not marry Rachel." The strongest proposition chosen true according to the one option is incompatible given K with the strongest chosen true according to the other.

Suppose Laban gives Jacob a further option: letting Laban decide which of the daughters he will marry. If Jacob "cannot make up his mind," he might be willing to follow this option;

but the strongest proposition chosen true in choosing this option is inconsistent with the strongest proposition chosen true in choosing to marry Rachel (or in choosing to marry Leah) and, hence, cannot be said to be weaker. Ranking options with respect to the logical strength (relative to K) of strongest propositions chosen true will not do.

Remember, however, that this method of ranking options with respect to strength is not followed in inferential expansion. Why should it be followed elsewhere? How, if at all, may we extend the notion of ordering options with respect to strength to the noncognitive case?

Keep in mind that the principle of P -admissibility is a "test" for the admissibility of feasible options to be administered once X 's knowledge, credal state, and values have been fully exploited in an evaluation of feasible options with respect to expected utility. To identify P -admissible options, new features of options are taken into account that were hitherto not emphasized.

Thus, in ranking expansion strategies with respect to strength, certain interesting features are abstracted from X 's evaluation of the outcomes of the several cognitive options with respect to epistemic utility. The M -function induces a total ordering of the potential answers with respect to informational value. This M -function agrees with the quasi ordering with respect to logical strength I recommend using. The weaker the potential answer, the higher the M -value. Thus, the ranking of options for purposes of testing for P -admissibility depends upon selecting features of the goals and values employed in evaluating the feasible options with respect to expected utility. It involves some aspect which, from the point of view of the agent, is or should be salient.

In constructing epistemic utility functions, I have recommended that in inferential expansion the proximate cognitive goals meet certain conditions. In so doing, I am advocating a certain vision of what the aims of such inquiry should be like. I have no proof of my thesis. The best that can be done is to explore the ramifications of this view to find out whether its consequences are acceptable in an account of good scientific inquiry.

In the same spirit, I contend that in evaluating options for P -admissibility, cognitive options should be ranked with re-

spect to strength in the manner indicated. This claim is another ingredient in my view of what the proximate aims of efforts at expansion should be. It is not relevant when attention is focused merely on considerations of expected utility. Nor is it an element in the account of rational choice I am proposing. It is no more a feature of my account of rational choice than my claim that in efforts at expansion scientists should seek error-free information.

When other types of decision making are considered with different goals and values, there may be no standard principle for selecting an aspect of the goals and values to be used in inducing a quasi ordering of options for purposes of testing *P*-admissibility. The specification of a quasi ordering can be regarded as itself an addendum to *X*'s goals and values pertinent to the decision problem under consideration.†

This approach does not trivialize the principle of *P*-admissibility. Even in cases where *X* counts all feasible options as noncomparable with respect to strength—as, I suspect, is the case in many practical decision problems—the principle remains applicable. It implies, under such conditions, that all and only *E*-admissible options are *P*-admissible.

On the other hand, in those cognitive decision problems where the method for ranking options with respect to strength is clearly indicated (which should be the case in contraction as well as inferential expansion), *P*-admissibility prescribes a different sort of result. The set of *P*-admissible options will not, in general, coincide with the set of *E*-admissible options. And similar results may obtain should some nontrivial way of ranking options with respect to strength be endorsed.

Suppose, in some practical decision problem, that *X* cannot make up his mind how to rank options with respect to strength. *X* may then employ the quasi ordering that stipulates that one

† Formally the proposal is that the partial ordering is relative to a partition *W*. For each feasible option, we can identify the strongest disjunction of elements of *W* to be added to the corpus if the agent elects that option. In inferential expansion, *W* is the ultimate partition *U*. Notice also that in inferential expansion there is one and only one option associated with each distinct proposition expressible as a disjunction of elements of *U*. This is not true for all decision problems. For each feasible option, there will be at least and at most one disjunction in *W*; but several feasible options may be associated with the same disjunction. Furthermore, there may be disjunctions which are not associated with any feasible options.

option is at least as strong as another if and only if all methods of ranking entertained by *X* agree. Otherwise the options are noncomparable.

Thus, the principle of *P*-admissibility does what it is designed to do. It is a principle applicable to all decision problems which stands in favor of suspension of judgment in those cases where considerations of expected utility fail to yield a definite recommendation. It expresses this view quite adequately in the context of inferential expansion and is formulated in such a manner that differences between cognitive and practical decision making are due to the differences between characteristic features of the aims of cognitive decision making and practical aims rather than to differences in principles of rational choice.

7.1 Decision Making Under Uncertainty

Is there a test for admissibility additional to the tests for E -admissibility and P -admissibility?

The question becomes important for applications to those practical decision problems where options are noncomparable with respect to strength. The test for P -admissibility is then ineffective for the purpose of weeding out E -admissible options. If, in addition, the agent's credal state is extremely indeterminate, the test for E -admissibility will also be nearly useless for the purpose of rendering a verdict.

The worst circumstance arises when X 's credal state is maximally indeterminate with respect to the decision problem under consideration. Each Q -function in B defines an n -tuple of conditional Q -distributions of the form $Q(h_k; A_i)$ over the h_k 's in the set U of states of nature (relative to the given classical representation of the decision problem) where there is one component in the n -tuple for each $A_i \in \mathcal{A}$. In a state of maximal indeterminacy, all such n -tuples consonant with X 's knowledge and the principles of inductive logic are permissible.

If the principles of inductive logic are restricted to credal coherence and direct inference as objectivists maintain, situations may arise where all n -tuples of Q -distributions satisfying the requirements of the calculus of probability relative to K count as permissible. Virtually all feasible options would be E -admissible. The only exceptions would be options so strongly dominated by others that, no matter which permissible u -function is considered, the best possible outcome for that option is worse than the worst possible outcome for some other option. All other dominated options and all undominated options would be E -admissible.

In a somewhat more determinate credal state, information concerning the option chosen is credally irrelevant for the hypotheses in U . Components of any given n -tuple of condi-

tional distributions are all identical with each other and with the corresponding unconditional Q -distribution over U . Yet, enormous indeterminacy remains if every unconditional Q -distribution over U obeying the calculus of probabilities is permissible. No option weakly dominated by any other will be E -admissible. All other options will be E -admissible. The increase in credal determinacy due to commitment to credal irrelevance of options for states does not carry very far.

In their classical compendium, R. D. Luce and H. Raiffa characterize decision making under risk as decision making where the probabilities of outcomes of options are given. In decision making under uncertainty, "the probabilities of these outcomes are completely unknown or are not even meaningful."¹

One of the advantages of the method proposed here for the representation of credal states and utilities is that it permits clearer and more sophisticated discriminations than those involved in this crude classification. In practice, Luce and Raiffa appear to restrict decision making under risk to cases where credal states are strictly Bayesian and options are credally irrelevant for states. I suspect that decision making under uncertainty was intended to cover cases corresponding to situations where maximal credal indeterminacy obtains, subject to the constraint that credal irrelevance also holds. This is not at all clear, however; and perhaps cases of maximal indeterminacy where credal irrelevance no longer obtains are also to be covered.

The important fact is that many authors for diverse reasons have attempted to formulate criteria for decision making under uncertainty. Among the criteria which have been proposed are minimax regret, optimism-pessimism criteria, maximin, and leximin.[†] All of these criteria rule out strongly dominated

[†] Luce and Raiffa, in *Games and Decisions* (New York: Wiley, 1958), ch. 13, discuss minimax regret as minimax risk. The criterion is attributed to L. J. Savage, who minimizes his own originality in *The Foundations of Statistics* (New York: Wiley, 1954), p. 170. Savage calls the rule minimax loss. But loss is often used as the negative of utility and, in this form, it is the minimax principle used by A. Wald in his *On the Principles of Statistical Inference* (Notre Dame, Ind.: University of Notre Dame Press, 1942). This latter principle is equivalent to the maximin principle discussed and so called by Luce and Raiffa, and it is under this name that I shall discuss it. Luce and Raiffa attribute the optimism-pessimism criterion to L. Hurwicz due to his

options—and leximin rules out weakly dominated options as well—as will become apparent when we formulate them. Thus, whether we employ them in contexts of decision making under uncertainty construed as presupposing or not presupposing credal irrelevance makes relatively little difference. In either case, we can construe these criteria as applying to all those options which have survived the tests of *E*-admissibility and *P*-admissibility—which, in effect, is to say that the criteria apply to virtually all of the feasible options.

Yet, there is a subtle but important difference between this way of approaching the matter and the way students of decision making under uncertainty customarily perceive the problem. No matter which criterion is employed, it is applied to the appraisal of all the feasible options rather than to that subset consisting of *E*-admissible (and hence, under the conditions assumed to prevail, *P*-admissible) options.

In practice this view makes little difference except for the evaluation of dominated options. But the conventional understanding of these criteria prevents recognition of the opportunities available for applying these criteria in contexts where there is more credal determinacy.

If, as proposed here, principles of choice designed for decision making under uncertainty are formulated as criteria for evaluating *P*-admissible options with respect to admissibility rather than for evaluating feasible options with respect to admissibility, it becomes apparent that the domain of applicability of these criteria may be extended beyond contexts of decision making under uncertainty.

As just indicated, however, there are many different criteria for decision making under uncertainty to consider. Those mentioned previously by no means exhaust the list of those which have been proposed. But they are the major contenders. We must somehow obtain some sort of appraisal of their merits.

I shall adopt a generalized version of a maximin criterion—

paper *Optimality Criteria for Decision Making under Ignorance* (Cowles Commission discussion paper, Statistics, No. 370, 1951). However, a more general approach of this sort had already been published by G. L. S. Shackle in *Expectation in Economics* (Cambridge: Cambridge University Press, 1949). The lexicographical maximin or leximin criterion is discussed in connection with welfare orderings in A. K. Sen, *Collective Choice and Social Welfare* (San Francisco: Holden-Day, 1970), p. 138.

although I am quite prepared to adopt leximin instead. In section 7.3, I shall offer some considerations in favor of choosing maximin or leximin. Before, however, doing so, it will be useful to furnish more explicit characterizations of all the criteria mentioned: minimax regret, optimism-pessimism, maximin, and leximin.

7.2 Regret, Hope, and Security

Given the states of nature in *U* according to some classical representation of a decision problem and given a permissible *u*-function in *G*, let u_k^* be the maximum value the *u*-function assigns an o_{ik} for fixed $h_k \in U$ and for all A_i 's in \mathcal{A} .

The regret $r(A_i; h_k)$ in implementing A_i conditional on h_k being true is equal to $u_k^* - u(o_{ik})$. The regret level $r(A_i)$ for A_i is the maximum value for $r(A_i; h_k)$ among all h_k 's in *U*.

The function $r(A_i)$ is determined by the states in *U* and the *u*-function in *G* and induces a complete ordering of the feasible options with respect to regret relative to the given *u*-function.

A feasible option is *regret optimal* relative to *u* if and only if its regret level is a minimum among the *P*-admissible options relative to *u*. Alternatively, it is a *minimax regret solution* relative to the *P*-admissible options.

An option is *regret admissible* if and only if it is regret optimal for some permissible *u*-function in *G* among the *P*-admissible options.

In discussions of decision making under uncertainty, it is often assumed that the decision maker has a set *G* of *u*-functions such that all members of *G* are linear transformations of one another. In effect, the regret admissible options are all minimax regret solutions relative to the same *u*-function. When *G* is allowed to violate the condition of uniqueness under linear transformation, minimax regret solutions relative to one *u*-function may not be minimax regret solutions relative to another and this accounts for the complication in the characterization of regret admissibility.

The other point to notice is that regret admissible options are characterized as minimax regret solutions among the set of all *P*-admissible options and not among the feasible options. In this way, the minimax regret criterion is rendered applicable to all contexts of decision making and not merely decision making under uncertainty.

In defining u_k^* , we might take it to be the maximum value

among all P -admissible options rather than all feasible options. It will give substantially the same results in contexts of decision making under uncertainty; but in more determinate situations, there can be a variation in the prescriptions made.

The *hope level* $h(A_i)$ relative to a u -function in G is the maximum value for $u(o_{ik})$ among all h_k 's in U . The *security level* relative to the same u -function is the minimum value $s(A_i)$ for $u(o_{ik})$ among the h_k 's in U .

For each option, the *optimism-pessimism* or *focal pair* consists of the hope level and the security level. G. L. S. Shackle, the originator of this criterion, introduced what one might call an adjusted optimism-pessimism pair or what he called standardized focus values.² I shall not introduce this complication here.

Let us suppose that we have some principles for ranking the various optimism-pessimism pairs which satisfies a Pareto condition. To fix ideas, we shall consider a ranking according to the weighted average $\alpha h(A_i) + (1 - \alpha)s(A_i)$ for $0 \leq \alpha \leq 1$. The ranking need not be of this form but could be so.

Relative to such a ranking using some permissible utility function in G , options among the P -admissible can be identified as *optimism-pessimism optimal* and *optimism-pessimism admissible*; OP -admissible options are those that are OP -optimal relative to some permissible u -function.

Strictly speaking, there are as many different notions of OP -admissibility as there are criteria for ranking optimism-pessimism pairs. Thus, for each choice of a value for the index α a different criterion is constructed.

When $\alpha = 0$, the criterion which emerges is a maximin criterion. Options are ranked relative to u with respect to security, and a definition of *security optimality* (S -optimality) among P -admissible options relative to u is easily constructed. An option is *S-admissible* if and only if it is S -optimal relative to some u -function—i.e., if it is a maximin solution for some u -function among the P -admissible options.

There may be many maximin solutions relative to a given u -function. One can rank these according to a so-called *lexicographical maximin* or *leximin* ranking. This compares options with the same security levels by considering the second worst possible outcomes and, if these are equal, looks at the third worst and so on until no more comparisons are available

or the tie is broken. *Leximin optimality* relative to u can then be defined and an option becomes *lex-admissible* if and only if it is lex-optimal among P -admissible options relative to some permissible u -function in G .

Regardless of whether one countenances regret admissibility, OP -admissibility, S -admissibility or lex-admissibility as necessary and sufficient for admissibility, in situations where the set of E -admissible and P -admissible options coincide, the options which become admissible are E -admissible options. They are, in this sense, Bayes solutions.

Some authors conclude from this fact that in choosing an option according to one of these criteria one is acting as if one's credal state is one where the uniquely permissible Q -function renders that option optimal with respect to expected utility.

For example, it is often alleged that maximin is a pessimistic procedure. The agent who uses this criterion is proceeding as if nature is against him.

It is then objected that such pessimism is unwarranted and, indeed, blatantly so; for it presupposes that any change in the agent's values will lead to an adjustment in nature's strategy so as to thwart the agent.

The interpretation of maximin involved in this objection is not the interpretation of maximin I am using.†

Maximin is a criterion for evaluating those options that have survived tests for E -admissibility and P -admissibility. Considerations of expected utility have failed to decide between these options. This is often due to credal indeterminacy. According to some permissible Q -functions in the credal state, the maximin solution is optimal with respect to expected utility. Had one of these Q -functions been uniquely permissible, that option would have been uniquely E -admissible. Given the agent's

† In *The Foundations of Statistics* (New York: Wiley, 1954), p. 181, Savage claims that "ultrapessimism" is unfairly charged to the minimax rule in decision theory by confusing it with the use of minimax in two-person zero-sum games where minimax loss means minimax the negative of income (i.e., maximin utility). Confusion is avoided, so Savage thinks, by interpreting minimax in decision theory as minimax regret. At the same time, so is the charge of ultrapessimism. It appears that Savage believes that a decision theory based on minimaxing negative income or utility is an ultrapessimistic theory. See also J. Harsanyi, "Can the Maximin Principle Serve as a Basis for Morality? A Critique of John Rawls's Theory," *Essays on Ethics, Social Behavior, and Scientific Explanation*, Dordrecht: Reidel, 1976, pp. 39-40.

values, that strictly Bayesian credal state would have been a pessimistic one.

But X does not endorse such a pessimistic credal state. Had he done so, invoking maximin would have been unnecessary. Considerations of expected utility alone would justify choosing what we are calling the maximin solution. It is precisely when the agent is neither committed to nor opposed to pessimism that he must go beyond expected utility and invoke other considerations in making a decision. But the decision made is to choose the maximin solution—not the pessimistic credal state. X should remain in suspense between the pessimistic Q -functions that render the maximin solution optimal with respect to expected utility and those Q -functions that are permissible according to his credal state but that do not render the maximin solution optimal. Considerations of security and of utility alone decide in favor of the maximin solution.

Strict Bayesians, of course, deny that there are any genuine contexts of decision making under uncertainty or of even intermediate credal indeterminacy. Hence, the problems we are dealing with ought not to be of concern to them. Many authors of this persuasion are prone to look upon advocates of maximin or one of the other criteria cited here as either worried or wishful thinkers who let their credal states (taken to be strictly Bayesian) be modified by changes in their values. Ingenious arguments are adduced to show that maximinners are closet Bayesians who lapse into incoherence because of their wishful thinking. But all such arguments (as well as similar arguments against the rival approaches described here) beg the question by interpreting the criteria of choice not as conditions on admissibility to be applied after considerations of expected utility (and other relevant factors such as P -admissibility) have failed to render a verdict but as covert ways of invoking considerations of expected utility.

We shall have occasion to explore a few examples of this sort of question-begging subsequently.† The question which needs to be faced now is whether there is any way of arbitrating between the alternative criteria mentioned in this section.

† Harsanyi, *ibid.*, pp. 39-40 and 46-47, furnishes ample illustration of a Bayesian dogmatism which refuses to countenance situations where credal states are indeterminate.

7.3 Spreads in the Odds

In the summer of 1973, I read in the newspapers that English bookmakers were prepared to bet 2 to 1 that Nixon would not be impeached and to take the other side of the gamble at odds of 1 to 3. They were not, however, prepared to take bets of the first kind at odds longer than 2 to 1 nor bets of the second kind nearer to 1 to 2.

Bookmakers are notoriously acquisitive men interested in a good rate of return on their money and this, no doubt, accounts in part for the spread in the odds. However, the spread is considerable and bookmakers tend to be more cautious than their customers. It is not at all implausible to suppose that their behavior at the time revealed an indeterminacy in their evaluations with respect to credal probability of the hypothesis that Nixon would be impeached.

To explain how spreads in the odds might arise due to credal indeterminacy, consider the following artificial gambling situation:

Case 1: h_A asserts that the chance of coin a landing heads on a toss is .4 and h_B asserts that the chance of a landing heads is .6. X knows that $h_A \vee h_B$ but does not know which of these alternatives is true. X also knows that a is to be tossed and will land heads (e_H) or tails ($\sim e_H$).

Suppose that X is offered a gamble on the outcome of the toss where he receives S dollars if e_H is true and nothing otherwise. X , so it shall be assumed, has neither taste nor aversion for gambling per se and, for such small sums of money, the utility of gains or losses in money is linear in the amount of the gain or loss. To all intents and purposes, X 's goals and values may be represented by any linear transformation of his monetary payoff function.

X is asked how much he is willing to pay for the gamble. If he pays P dollars, then if e_H turns out to be true, X receives $S - P$ dollars; and if false he receives $-P$ dollars. Clearly a gamble where P is reduced is favored over one where it is increased. But the issue is how high P becomes before X refuses the gamble and, as a consequence, receives 0 dollars regardless of the outcome of the toss.

The question presupposes that there is, indeed, a threshold

price beyond which X will refuse the purchase of the gamble when the stake S is fixed.

Suppose X 's credal state were strictly Bayesian so that $Q(e_H) = r$. We can, in point of fact, determine that r must be some number greater than or equal to .4 and less than or equal to .6.

In virtue of the principle of direct inference (whose formulation will come in later chapters) and given X 's knowledge, every Q -function in his credal state should satisfy the requirements that $Q(e_H; h_{.4}) = .4$ and $Q(e_H; h_{.6}) = .6$. Credal coherence requires, therefore, that $Q(e_H) = .4Q(h_{.4}) + .6Q(h_{.6})$, where $Q(h_{.6}) = 1 - Q(h_{.4})$. Hence, given a value $x = Q(h_{.4})$ somewhere between 0 and 1, a value for $Q(e_H)$ somewhere between .4 and .6 will be determined. If $Q(e_H) = r$ for all permissible Q -functions in the credal state, then all permissible Q -functions assign $h_{.4}$ the same numerical value and, for the purposes of the problem, the situation is strictly Bayesian.

The threshold value for P can then be determined. It is the price for which the expected utility of the gamble utilizing $r = Q(e_H)$ is equal to the expected utility of refusing the gamble, i.e., 0. That is to say, $r(S - P) - (1 - r)P = 0$. Hence, $P = rS$.

The ratio $P/S = r$, where P is the threshold price, is called a fair betting quotient for the gamble. This ratio should be the same regardless of variations in the magnitude of the stake S and regardless of its sign. If $S > 0$, the fair odds are $P/(S - P)$. If S is negative, the fair odds are $(S - P)/P$ —i.e., the reciprocal of the fair odds when $S > 0$. The fair odds when S is negative are odds against e_H .

These are the consequences in the strict Bayesian case. There can be no spread in the odds. If, as is the case in real life, there are spreads in the odds and the credal state is strictly Bayesian, the behavior is rational only if utility is nonlinear in the monetary payoff (due to taste for or aversion to gambling or to a lust for profits of a fixed minimum size or the like). In our hypothetical case 1 we can stipulate that such factors are to be ignored.

Following C. A. B. Smith,³ I contend that in cases like case 1, even in real life X 's credal state will not always or even often be strictly Bayesian. It will allow many Q -values to be

permissible for $Q(h_{.4})$ and, as a consequence, the set of permissible Q -values for e_H will form a subinterval of the interval from .4 to .6. (All the points in the subinterval, except perhaps for endpoints, will be permissible due to the convexity condition.) To simplify the discussion, suppose that the interval is the entire interval from .4 to .6.

Smith's thesis is that in this kind of situation there will continue to be a threshold price for stake S where $S > 0$ and a threshold price for stake S' (of the same magnitude) where $S' > 0$. However, P/S will be less than P'/S' . That is to say, the odds at which X is willing to bet on e_H are not the reciprocal of the odds at which he is willing to bet against. There is a spread in the odds, as in the case of the London bookmakers, due here to indeterminacy in credal state.

Indeed, Smith assumes that the threshold betting quotient for $S > 0$ —i.e., the lub of betting quotients for positive stakes at which X will accept rather than refuse the gamble—should equal .4. The threshold when $S < 0$ —i.e., the glb of betting quotients at negative stakes at which X will accept rather than refuse the gamble—should equal .6.

In general, Smith's contention is that when $S > 0$, the threshold is the endpoint of the interval of permissible Q -values for the hypothesis in question to the left or the "lower pignic probability"; and when $S < 0$, the threshold is the endpoint to the right or the "upper pignic probability."

I contend that, in the face of credal indeterminacy of the sort described, rational agents should appraise such gambles in a manner which allows for such a spread in the odds. I do not know how to prove this to be so; and if someone is unconvinced I can offer nothing to convince him. However, one small advantage of making this assumption is that, given that the agent is rational and that utility is linear with respect to the monetary payoff, one might entertain the devising of circumstances which reveal whether his credal state for e_H is numerically precise or indeterminate.

According to the account of rational choice being proposed here, how does the case 1 situation appear with indeterminate credal state? When the betting quotient P/S is less than .4 and S is positive, the expected utility of accepting the gamble is positive regardless of the Q -function, permissible according to the credal state, that is used to compute expected utility.

Accepting the gamble is uniquely E -admissible, uniquely P -admissible, and, hence, uniquely admissible. If $P/S > .6$, refusing the gamble is uniquely admissible by similar reasoning.

If P/S falls in the interval from .4 to .6, both accepting the gamble and refusing it are E -admissible and, hence, P -admissible.

According to Smith's proposal, if .4 is the lowest permissible Q -value for e_H and .6 the highest such value, gambles for such "medial" betting quotients should be rejected and the threshold betting quotient should be .4.

By similar reasoning if $S < 0$, then the gamble is uniquely E -admissible for betting quotients greater than .6 and refusing the gamble uniquely admissible for betting quotients less than .4. Once more Smith's prescription recommends rejecting gambles for medial betting quotients.

Thus in hypothetical gambling situations such as case 1, wherein accepting a gamble for a given stake and betting quotient is compared with rejecting it, Smith's proposal does give a condition for determining which E -admissible (and, hence, P -admissible) options are admissible.

On the other hand, the question remains open as to how to proceed in general. Suppose there are three P -admissible options or more than three. Or suppose there are two P -admissible options but different payoffs for three or more states of nature.

In such cases, I think Smith regards all E -admissible and, hence, P -admissible options to be admissible. But I am not sure. If, however, this is his view, it stands in contradiction to his treatment of pairwise choices of the sort exemplified in case 1.

A general approach can be obtained, however, by considering the various criteria for decision making under uncertainty as generalized previously. We can ask, for each of them, whether they allow for spreads in the odds in case 1 situations. If they do not, there are at least some slender grounds for dismissing them.

I propose to use Smith's approach to case 1 predicaments as a condition of adequacy for evaluating these criteria; but I do so very tentatively.

Of the four criteria canvassed, only the criteria of S -admissibility and lex-admissibility may be used as conditions for

admissibility while permitting spreads in the odds in case 1 situations in the manner described by Smith.

That these two criteria pass is apparent. There are only two possible consequences with distinct payoffs if the gamble is chosen and one if the gamble is refused. Consequently, when these options are E -admissible, they can qualify as S -admissible if and only if they are lex-admissible. In these cases, the one criterion yields the same results as the other.

Furthermore, refusing the gamble always bears higher security than accepting it (as long as $S > P$). Consequently, when both options are E -admissible (and, hence, P -admissible), the criterion of S -admissibility favors refusing the gamble. Hence, for medial odds, the gamble should be refused just as Smith requires.

If regret-admissibility or OP -admissibility are used as necessary and sufficient for admissibility, results conflicting with Smith's prescriptions can occur.

Consider regret admissibility. In case 1 when $S > 0$, accepting the gamble has regret 0 when e_H is true and regret P when e_H is false. Refusing the gamble has regrets $S - P$ and 0, respectively. Refusing the gamble has lower regret level than accepting it if and only if $S - P < P$, i.e., if and only if $P/S > .5$. And if S is negative, accepting the gamble has regret $P - S$ and 0 for the two states whereas refusing has regrets 0 and $-P$. Refusal has lower regret if and only if $-P < P - S$, i.e., if and only if $P/S < .5$.

Consequently, there can be no spread in the odds in our case 1 predicament. The threshold betting quotient is the same whether S is positive or negative and it is different from the lower and upper permissible Q -values. If the interval in other variants of case 1 fails to cover the value .5, then if the upper credal or pignic probability is less than .5, the threshold is at that upper value. If the lower credal probability is greater than .5, the threshold betting quotient is at that value. These consequences clearly conflict with Smith's prescription.

Using OP -admissibility sometimes allows for a spread in the odds. Recall, however, that Smith requires that the threshold for betting quotients when S is positive is to equal the lower probability and the threshold when S is negative is to equal the upper probability. As a consequence, the former threshold is less than or equal to the latter.

But if α is greater than .5 this cannot be the case; for when S is positive and P/S is medial, the gamble is accepted if and only if $P/S < \alpha$. When S is negative and P/S is medial, the gamble is accepted if and only if $P/S < (1 - \alpha)$. When α is greater than .5, $(1 - \alpha) < \alpha$.

So let us restrict α to values less than .5. We obtain a spread in odds but the thresholds will not coincide with the lower and upper credal probabilities unless α is less than or equal to .4. By changing the examples so that the lower probabilities are arbitrarily close to 0 and the uppers arbitrarily close to 1, the value of α is forced to 0. But when $\alpha = 0$, OP -admissibility coincides with S -admissibility.

Let us then sum up the argument. Four well-known criteria for decision making under uncertainty have been reformulated so that they are applicable not only in contexts of maximal credal indeterminacy but in contexts of intermediate credal indeterminacy as well. They have also been modified so that sets of utility functions may represent payoff structures even though they are not unique up to a linear transformation.

These four criteria have then been applied to a hypothetical gambling situation where partial credal indeterminacy obtains. Under the special circumstances of that kind of decision problem, Smith's view that gambles at medial odds should be refused is adopted as a standard for evaluating the applicability of our four criteria to such situations. Of the four criteria, minimax risk and OP -admissibility fail and maximin (or S -admissibility) and leximin (or lex-admissibility) pass.

On this tenuous basis, I propose to endorse either S -admissibility or lex-admissibility as necessary and sufficient for admissibility. For simplicity's sake and because it has been more widely advocated, I shall adopt S -admissibility here. Most of what will be said about S -admissibility applies *mutatis mutandis* to lex-admissibility so that this decision, though arbitrary, makes no significant difference to the development of the themes to be discussed.

7.4 On Appraising Security

A. Wald pioneered in the use of maximin criteria in statistical decision making. He thought that through the use of a statistical decision theory based on maximin he could systematize and generalize the approach of J. Neyman and E. S. Pearson to the handling of statistical data.⁴

As I understand them, Neyman, Pearson, and Wald were objectivist necessitarians (see chapter 4). They recognized that objectivist necessitarians are faced with a serious challenge to explain how data can be useful in deliberation and inquiry and, in particular, in experimentation. This challenge will be explained in chapter 14; the Neyman-Pearson-Wald response will be discussed in chapter 17. Wald undertook to make a version of the maximin criterion the centerpiece of his own response. As a preliminary to those later discussions, some comparison of Wald's version of maximin and the one used here to construct a criterion of S -admissibility is in order.

The criterion of S -admissibility was built on the basis of the notion of a security level for an option relative to a permissible u -function. Given that notion, P -admissible options could then be ranked with respect to security relative to the u -function and an S -admissible option understood to be one which is optimal with respect to security among P -admissible options according to the security ranking relative to some permissible u -function in G .†

Thus, the applicability of the criterion of S -admissibility depends on how security levels are determined relative to u -functions. The previous discussion has glossed over an important ambiguity in the instructions for doing this which must now be faced. It lies at the heart of the issue which Wald's version of maximin raises.

Consider the following hypothetical decision problem:

Case 2: $h_{.4}$, $h_{.6}$, and e_H are as defined previously. X has a choice between three options: G_1 is a gamble with a stake of one dollar and a price of \$0.45 on the coin landing heads on a toss. G_2 is a gamble with a stake of \$2.00 and a price of \$0.80 on the coin landing tails. G_3 is the option of refusing the gamble with neither loss nor gain. X must choose exactly one of these options.

The payoff matrix is shown in table 7.1.

Assume that options are credally irrelevant to states. Direct inference requires that $Q(e_H; h_{.4}) = Q(\sim h_H; h_{.6}) = .4$ and

† Wald does not appear to have restricted his criterion to P -admissible or even E -admissible options but there is some small indication that he might have been prepared to do so (*On the Principle of Statistical Inference*, Notre Dame, Ind.: University of Notre Dame Press, 1942, p. 44).

Table 7.1 Payoff Matrix

	$h_{.4} \& e_H$	$h_{.4} \& \sim e_H$	$h_{.6} \& e_H$	$h_{.6} \& \sim e_H$
G_1	.55	-.45	.55	-.45
G_2	-.80	1.20	-.80	1.20
G_3	0.00	0.00	0.00	0.00

$Q(e_H; h_{.6}) = Q(\sim e_H; h_{.4}) = .6$. Let r be the Q -value for $h_{.4}$ and $(1 - r)$ for $h_{.6}$. Then the Q -values for the various states of nature become $.4r$, $.6r$, $.6(1 - r)$, and $.4(1 - r)$, respectively. The expected utilities for the three options relative to the Q -function assigning $h_{.4}$ the unconditional value r are

$$E(G_1) = .22r - .27r + .33(1 - r) - .18(1 - r),$$

$$E(G_2) = -.32r + .72r - .48(1 - r) + .48(1 - r),$$

$$E(G_3) = 0.$$

To identify which of the feasible options are E -admissible requires a specification of the range of permissible Q -values for $h_{.4}$ —i.e., the range of values which r may take. Let us, for the moment, assume it is the maximum range from 0 to 1. If $r = 0$, then $E(G_1) = .15$ and $E(G_2) = E(G_3) = 0$. As r increases, the expected utility of G_1 declines to a minimum of $-.05$ when $r = 1$. The expected utility of G_2 increases from 0 to $.4$. The expected utility of G_3 remains 0.

Consequently, among these three options, G_3 fails to be optimal relative to any permissible Q -function and, hence, is E -inadmissible. The other two options are both E -admissible and presumably P -admissible. Considerations of security need to be invoked to arbitrate between them.

Notice that, following on the way possible consequences are partitioned and payoffs specified in table 7.1, G_3 bears the highest security level, G_1 is next, and G_2 bears the lowest security level. But G_3 is ruled out of court because of considerations of expected utility. We need to consider only G_1 and G_2 . G_1 seems uniquely S -admissible and, hence, admissible. It should be chosen.

But there are other ways to look at the matter. According to table 7.1, four hypotheses as to the outcome of an option are specified. Thus, for option G_1 these hypotheses are that G_1 is implemented while $h_{.4} \& e_H$ is true, while $h_{.4} \& \sim e_H$ is

Table 7.2

	e_H	$\sim e_H$
G_1	.55	-.45
G_2	-.80	1.20
G_3	0.00	0.00

true, while $h_{.6} \& e_H$ is true, and while $h_{.6} \& \sim e_H$ is true. Similar characterizations may be offered for G_2 and G_3 .

Observe, however, that e_H is equivalent given the background knowledge K to $h_{.4} \& e_H \vee h_{.6} \& e_H$. Given the implementation of any one of the options, the payoff is the same regardless of which of the two disjuncts in e_H is true. Hence, no matter which permissible system of Q -values is adopted, the utility assigned implementing option G_i while either $h_{.4} \& e_H$ or $h_{.6} \& e_H$ is true is the same. If the option is G_1 , it is $.55$, G_2 yields $-.80$ and G_3 gives 0. The u -value for implementing G_i while e_H is true is *credally insensitive* in this sense.

The same applies to $\sim e_H$. We can rewrite our payoff matrix as shown in table 7.2.

The permissible Q -values for e_H range from $.4$ to $.6$ as r ranges from 0 to 1. Obviously the evaluations with respect to E -admissibility remain as before. So do the evaluations with respect to S -admissibility.

Clearly what we have done here is derived a new system of states of nature and system of possible consequences for each feasible option by taking disjunctions of possible consequences in the original "partitions" to form a new system of partitions. The credal insensitivity to which reference was made previously enabled us to determine u -values for each possible consequence in the new system independent of a selection of a Q -function from the set permissible according to the credal state.

Observe, however, that there is another way to coarsen these partitions. $h_{.4}$ is the disjunction of $h_{.4} \& e_H$ and $h_{.6} \& \sim e_H$, and $h_{.6}$ may be obtained in a similar manner.

Now given the truth of $h_{.4}$, the ultimate payoff for any option (except G_3 , which remains 0 throughout) depends on whether e_H is true or false. Nonetheless, a definite u -value can be

attributed to $h_{.4}$ given implementation of an option. For option G_1 , that expectation is

$$E(G_1; h_{.4}) = Q(e_H; h_{.4})(.55) - Q(\sim e_H; h_{.4})(.45).$$

This u -value for $h_{.4}$ would be credally sensitive were it not for the fact that the principle of direct inference, so we are assuming here, mandates that every permissible Q -function meet the condition that $Q(e_H; h_{.4}) = .4$. Hence, regardless of the permissible Q -function in the credal state being used, the utility assigned $h_{.4}$ when G_1 is implemented becomes $.22 - .27 = -.05$. Similar computations for $h_{.6}$ and for all options yields the payoff matrix shown in table 7.3.

Like the payoff structure shown in table 7.2, this payoff structure is obtained by coarsening the initial structure in a credally insensitive manner. Consequently, the resulting payoff structure has no effect upon how E -admissibility is to be evaluated. G_1 and G_2 come out E -admissible. G_3 does not. The representations, though different, are—so it seems—of the same decision problem. Yet, in table 7.3 the security level of G_2 is higher than G_1 —counter to the result according to the matrix of table 7.1 or 7.2.

Thus, we can see that formulations of a given decision problem which are equivalent as far as evaluations of expected utility are concerned can, nonetheless, yield distinct appraisals with respect to S -admissibility. And the difference is traceable to the way security levels are determined.

Wald and other students of statistical decision theory formulate decision problems so that the states of nature are statistical hypotheses such as $h_{.4}$ and $h_{.6}$ specifying what the chances of various outcomes of some sort of experiment bearing on the payoff will be. The utility assigned to choosing feasible option A_i when such a statistical hypothesis is true is the expected utility of A_i conditional on that hypothesis. This expected utility can be computed from the payoffs in utility after the process has been carried to completion for some u -function because direct inference (so, at any rate, it is assumed) determines the requisite conditional Q -values for the outcomes of experimentation given the statistical hypothesis. Thus, in case 2, $Q(e_H; h_{.4}) = .4$ and $Q(e_H; h_{.6}) = .6$ regardless of the Q -function being used to compute expectations. The payoff matrix of table 7.3 illustrates the procedure.

Table 7.3

	$h_{.4}$	$h_{.6}$
G_1	-.05	.15
G_2	.40	0.00
G_3	0.00	0.00

But payoff matrices 7.1 and 7.2 also may be used to determine security levels and, in our example and in many others, it can be shown that the options which become S -admissible will be different. I shall call the procedures for fixing security exemplified in these matrices *method (a)* and Wald's procedure *method (b)* for fixing security levels. In our case 2 example, there are other methods available as well; but it will suffice for present purposes if we consider these two.

Is there any right or wrong concerning the fixing of security levels according to method (a), method (b) or some other method?

I have taken for granted all along that an account of rational choice should not mandate what a rational agent ought to do in any given situation. The agent's knowledge, credal state and values need to be specified. Given this specification, I have proposed a criterion for evaluating feasible options with respect to E -admissibility. However, no stipulation concerning what X 's goals and values ought to be (except for minimal consistency requirements) has been imposed. To do more than this would be to enter into questions of morals and politics. I do believe that there is, in some suitably qualified sense, a right and a wrong about such matters; but the determination of right and wrong is often difficult and, to a large degree, criteria for rational choice may be proposed without settling the issues under dispute. Decision theorists ought not to be moralists, politicians, or guidance counselors.

What is obvious in the case of appraisals of E -admissibility ought to be obvious also in the case of appraisals of P -admissibility and S -admissibility. Thus, I contended that evaluations of feasible options with respect to strength to be used in assessing P -admissibility reflect the agent's goals and values and that decision theory should not legislate what these should be.

The same point applies to fixing security levels. How X selects security levels is a reflection of another aspect of X 's goals and values. Decision theorists should not legislate how this should be done, except to insist that it be done in a consistent and coherent manner.

My contention is, therefore, that insofar as Wald and others have advocated method (b) rather than method (a) or some other method for fixing security levels in determining S -admissibility, they have been engaged in moral and political argumentation and have gone beyond a discussion of rational decision making.

In a free society, there should be no objection to moral exhortation of this sort. However, intellectual clarity requires that men not be condemned as irrational because they refuse to fix security levels in one manner rather than another.

Thus, an investor on the stock market may not know how the Congress will decide on some taxation policy. If they decide one way, a given investment will have one expected value. If they decide in another, the investment will have another expected value. The Wald prescription is that the security level for the investment should be the smallest of these expectations. Method (a) prescribes attending to the worst that can happen when the investment is exchanged for money or some other commodity. The investor who looks to the ultimate payoffs in assessing security is no less rational than the agent who attends to the expectations conditional on the congressional decision. He has different values regarding security.

Suppose agent X is conflicted as to whether to use method (a) or method (b) or some other method for fixing security levels. How should he proceed? The spirit of the approach I have been advancing urges the agent to consider as S -admissible any option which comes out S -admissible relative to some method of fixing security levels taken seriously.

7.5 Mixed Options

The import of this conclusion for Wald's approach to statistical theory will be developed further in chapter 17. One important ramification can be mentioned here.

Suppose case 2 is modified by furnishing X with a fourth option G_4 . A spinner is spun with a .8 chance of falling in area 1 on a dial and a .2 chance of falling in area 2. If it falls in the

first area, X receives payoffs from the toss of the coin in accordance with G_1 ; if it falls in area 2, X receives payoffs in accordance with G_2 .

G_4 is a *mixture* of the pure options G_1 and G_2 . By direct inference we know that regardless of the permissible Q -function used to compute expected utilities, $E(G_4) = .8E(G_1) + .2E(G_2)$. Also $E(G_4; h_{.4}) = .8E(G_1; h_{.4}) + .2E(G_2; h_{.4}) = .36$. By similar calculation, $E(G_4; h_{.6}) = .12$.

This option is E -admissible because for $Q(h_{.4}) = r = .5$, $E(G_4) = .24$ whereas $E(G_1) = .05$ and $E(G_2) = .2$. On the other hand, G_1 and G_2 remain E -admissible. (Consider values of r near 1 and 0.) P -admissibility, once more, makes no difference. Considerations of security must decide.

If we follow Wald's method (b), the security level for G_4 will be $E(G_4; h_{.6}) = .12$ which is superior to that of the other two E -admissible options. Hence, the mixed option will be favored over the pure one on grounds of security.

On the other hand, if method (a) is employed, the mixed option G_4 has the same security level of $-.80$ as the pure option G_2 , which is inferior to the security level of G_1 . Not only does it fail to improve security levels, it should positively be avoided.

Notice that there is a third method for fixing security levels here. One could use Wald's method for the pure options and then fix the security level of the mixed option as the worst outcome of the two pure options (as method (a) requires). Once more, the mixed option furnishes no benefit.

Wald's method for fixing security levels is, of course, not original with him. In discussing two-person zero-sum games, J. Von Neumann and O. Morgenstern exploit mixed strategies to gain the existence of a solution to every two-person zero-sum game.⁵ Of course, the solution is gained by identifying the security levels of mixed strategies in a manner analogous to Wald's method.

However, if a decision maker facing such a game does not fix security levels in this manner, this does not show that he is irrational. It only demonstrates a limitation of the Von Neumann-Morgenstern theory of two-person zero-sum games.

8.1 Lack of Self- Awareness

That men fail to appeal explicitly to utility functions in ordinary deliberation is a familiar fact. Strict Bayesians know this and are rarely disturbed by it. Even when the attribution of strict Bayesian rationality is made an explanatory or prescriptive claim, strict Bayesians do not maintain that the agent explicitly or consciously identifies his utility function, his credal probability function, or his expected-utility function. The assertion that X 's values are representable by a utility function unique up to a linear transformation is a theoretical claim introduced for purposes of prediction or explanation of X 's choices and his answers as to what his choices would be in various hypothetical decision problems.¹

Of course, if strict Bayesian conditions are taken as prescriptions designed for policing credal states, goals and values, and choices in deliberation, the norms are applicable only insofar as X is able to apply them. Applying the norms may require becoming explicitly aware of the evaluations. But limitations of time, memory, computational facility, imagination, and emotional health impose serious bounds on X 's ability to identify a utility function unique up to a linear transformation representing his goals and values. If X should begin to make progress in making such identification, he may have already modified these values. Psychological inhibitions and defenses may engender self-deception. Or X may be conceptually confused. Such factors may place insuperable obstacles in the way of X 's identification of his goals and values. Unbeknownst to himself, X 's values and goals may be impeccable according to strict Bayesian standards; but unless he can identify the appropriate utility functions to a reasonably good degree of approximation, this will help him very little in determining which of alternative options he should choose within the particular context in which he finds himself.

Defenders of the Bayesian faith sometimes take this objec-

tion seriously. They recommend that X find out whatever he can about himself relevant to identifying his utility function. Even if the information he obtains is insufficient for singling out his utility function, it may be strong and accurate enough to show that some option is uniquely optimal with respect to expected utility and, hence, is uniquely admissible. Even when such definiteness is not to be obtained, some options may be eliminated on the grounds that they are clearly not optimal.²

This approach will seem unsatisfactory to advocates of the principle of maximizing expected utility or of the principle of S -admissibility I have adopted. When X is compelled to choose between options with partial knowledge of his own values, he cannot guarantee that he will maximize expected utility relative to a utility function correctly representing his system of values. At the very best, he can avoid choosing an option which he knows to be inadmissible according to those values.

One could try the following resolution of the difficulty. If X is ignorant as to which of rival u -functions truly represents his goals and values, he should have a credal state for the rival hypotheses as to what that u -function is. Given any particular hypothesis, an expected-utility function for the feasible options is defined. Multiply each such function by the probability that the u -function used is the correct representation of X 's goals and values and take the sum. This provides a new expected-utility function to be used in evaluating the options.³

Several objections are decisive against this view:

(1) In effect, X is using a utility function to compute expected utilities which is a weighted average of the utility functions between which he suspends judgment. The weights are the degrees of credence he assigns the rival hypotheses as to which u -function is the correct one. The weighted average would often, however, be different from any of the u -functions between which X suspends judgment. Hence, from X 's point of view, it is certain that the computed u -function does not represent his goals and values; hence, using it is a betrayal of these values.

(2) X 's lack of self-knowledge can extend to his credal state. If we follow the same procedure for the credal state which is suggested for the u -function, X should have a credal state over rival hypotheses as to what his correct Q -function is.⁴ The

new Q -function will differ, in many instances, from those between which X suspends judgment. Moreover, if X is in doubt as to what his credal state over his credal state is, he may have to move to a still higher level. Obviously we are at the brink of an infinite regress.⁵

(3) Suppose X suspends judgment as to whether u_1 or u_2 is the correct representation of his goals and values. Of course, if u_1 is the correct representation, so is any linear transformation thereof. The same holds for u_2 . Let r be the credal probability that u_1 is the correct utility function and $(1 - r)$ the credal probability that u_2 is the correct utility function. The proposal being canvassed suggests that $ru_1 + (1 - r)u_2$ be adopted as the utility function. But why should X not use $ru'_1 + (1 - r)u_2$ instead, where u'_1 is some linear transformation of u_1 ? If he does, the new weighted average will not be a linear transformation of the original one.

Clearly ordinary mortals are quite incapable of identifying their goals and values with sufficient clarity so as to represent them by means of convex sets of utility functions or by single utility functions. To this extent, in real life it will often not be feasible to implement strict Bayesian prescriptions or prescriptions based on S -admissibility.

Recall, however, that the same is true of the principle of deductive closure imposed on a corpus of knowledge. In section 1.5, I contended that X is committed to regarding as not possibly false (in the sense of serious possibility) all consequences of items in his corpus. It was not required that he live up to his commitments but only that he live up to his commitments insofar as he is able.

In the same spirit, we need not interpret strict Bayesians as flouting the principle that "ought" implies "can." They may regard systems of values and credal states as commitments that agents embrace at given times, and understand the conditions of rational valuation and credence as imposing constraints on such commitments. Bayesians may recognize that limitations of memory, computational facility, imagination, and emotional health may prevent complete identification of these commitments and obedience to them. They need not, therefore, tailor their views to what the best psychology and sociology tells us is feasible. To the contrary, when deductive closure cannot be obeyed because of failures of memory, com-

putational facility, or imagination, deductive closure should not be abandoned as a principle of rationality. Instead the advice of expert mathematicians is sought, computers are employed, and techniques for improving the memory and retrieving information are devised. When Bayesian conditions of rationality seem beyond our grasp, strict Bayesians can reasonably urge us to employ the available technology and psychotherapeutic techniques to enable us to live up to our commitments better than we have previously done.

No interesting system of norms can be applied under all circumstances. Provision must be made for excuses due to lack of feasibility. We should, therefore, recognize a distinction between principles of rationality regulating an agent's commitments and the suggestions which may be made when he cannot live up to them. That distinction parallels the familiar distinction drawn by statisticians between "exact" methods of estimation and "approximate" ones.

To be sure, any proposed system of principles of rational credence or valuation intended to be imposed on commitments is not entirely exempt from the demand that it demonstrate its applicability. What is needed is some indication that at least in principle methods can be devised for identifying commitments with the precision required for real-life application. Moreover, some effort should be made to establish that the practical applicability of the theory is not excessively limited or, if it is, to show that "approximate" methods or "suggestions" based on the principles regulating rational commitments can be concocted which render the theory relevant in real life.

Many efforts have been undertaken by strict Bayesians to show that their principles understood as constraints on commitments are at least in principle identifiable by methods of interrogation concerning hypothetical choice situations. And some Bayesians have focused attention on the problem of what an agent should do when he cannot live up to his commitments.⁶ Some of these matters are subject to controversy. However, I do not think a decisive objection can be leveled against strict Bayesians because ordinary agents are often incapable of numerical representations of their values or their credal states.

To my knowledge serious critics of the deductive-closure requirement do not complain of its applicability, but raise

other objections which purport to show (wrongly I think) that a rational agent should not commit himself to all deductive consequences of items in his corpus.⁷ The dispute is over what the commitments should be, and not whether the commitments proposed can always be satisfied. The same attitude should prevail in confronting strict Bayesian doctrine concerning rational valuation and rational credence.

A strictly Bayesian valuation of possible consequences of feasible options in a decision problem is representable by a utility function defined over these hypotheses which is unique up to a linear transformation. According to what has just been said, this utility function represents *X*'s commitments as to how hypotheses about consequences should be evaluated with respect to their success in promoting his values and goals.

I reject this strict Bayesian view of conditions which such commitments should satisfy. To repeat, my objection is not that ordinary agents are incapable of living up to such commitments. Rather it is that rational agents should not embrace such commitments in the first place. Even if *X* is so ideally situated that he is able to identify his commitments as to how possible consequences are to be appraised with respect to utility, he sometimes should not endorse commitments representable by a utility function unique up to a linear transformation. Although I do maintain that these commitments should satisfy the requirements of valuational consistency, closure under linear transformation, and valuational convexity cited in section 4.7, the additional strict Bayesian requirement of valuational uniqueness should not only not be imposed but in some contexts is positively to be violated.

8.2 Unresolved Conflict

Jacob is compelled to choose between marrying Leah (A_1), marrying Rachel (A_2), and marrying Zilpah (A_3). Jacob knows that if he chooses A_1 , he will be marrying an extremely ugly woman but superb cook (o_1). If he chooses A_2 , he will be marrying a divinely beautiful woman but poor cook (o_2). Finally if he chooses A_3 he will be marrying a woman of mediocre beauty and culinary ability (o_3).

Jacob faces a decision problem under certainty—i.e., where he is certain of what the consequences of each feasible option will be under a relevant description. The Bayesian principle of maximizing expected utility reduces in this case to selection

of that option or one of those options whose consequence bears maximum utility. To apply the principle Jacob need only have a valuation of consequences representable by a utility function over the o_i 's.

Indeed, we do not need even this much. Jacob need have only a rank ordering of the o_i 's with respect to his values.

If Jacob cared only for a wife who cooked well but not at all about beauty, he would rank o_1 over o_3 over o_2 . Marrying Leah would be uniquely *E*-admissible and admissible.

If Jacob cared only for a beautiful wife regardless of how she cooked, he would rank o_2 over o_3 over o_1 . Marrying Rachel would be uniquely *E*-admissible and admissible.

Suppose, however, that Jacob is seeking both a beauty and a good cook. Ideally he would want a wife at least as beautiful as Rachel and as good a cook as Leah. He does not have the opportunity to marry such a paragon. If he marries Leah, he sacrifices beauty for good cooking, and if he marries Rachel, he sacrifices cooking for beauty. If he marries Zilpah, he settles for mediocrity in both.

As long as Jacob is conflicted between these two desiderata and has failed to resolve the conflict, neither of the rankings described above represents his goals and values.

Indeed, there is no ranking in this case which represents his system of valuations. o_1 is neither better than, worse than, or equally as good as o_2 (or o_3) as far as Jacob is concerned. The three consequences are not comparable with respect to utility or value.

A fortiori there is no utility function unique up to linear transformation which represents Jacob's goals and values.

Bayesians are compelled to deny that Jacob's values are rational if they are of the sort described. But they cannot seriously deny that rational men can often be conflicted in their values as Jacob is conflicted between seeking a beauty and a good cook. Hence, Bayesians must deny that in cases like Jacob's, the agent's commitments are not representable by a utility function unique up to a linear transformation.

One method open to the Bayesians for doing this is to admit the presence of conflict but contend that it is resolved.

Suppose utility function u_1 represents the valuation of the o_i 's according to Jacob insofar as he focuses on his concern to marry a beautiful woman: $u_1(o_2) > u_1(o_3) > u_1(o_1)$. The

function u_2 represents the valuation with respect to cooking: $u_2(o_1) > u_2(o_3) > u_2(o_2)$.

Each of these utility functions taken by itself represents a strictly Bayesian goal where there is no conflict. The two goals are conflicting goals. But when taken separately the u -functions can be used to evaluate E -admissibility. If X were offered a lottery in which he was offered the prospect of marrying one of the three women with definite probabilities and was committed to a goal representable by u_1 , he could evaluate the goal in terms of expected utility using u_1 . If Y were offered a lottery where the payoffs were the same but the function u_2 represented his goals and values, u_2 could be used to compute expected utility.

Jacob's goals and values are conflicted because they are representable by a utility function u_α which is a function of u_1 and u_2 and, in particular, is equal to $\alpha u_1 + (1 - \alpha)u_2$, $0 \leq \alpha \leq 1$.

The function u_α is a resolution of the conflict between u_1 and u_2 . It involves a relative evaluation of the two desiderata represented by u_1 and u_2 . This relative evaluation can be represented by "weights" attached to the two utility functions. If Jacob attaches 0 weight to u_1 and weight of 1 to u_2 , he cares only for a good cook. Alternatively, assigning $\alpha = 1$ and $(1 - \alpha) = 0$ reveals Jacob as concerned only with beauty.

Strict Bayesians who follow this approach deny that Jacob, if he is rational, has a genuine conflict in values. When Jacob is alleged to have a conflict in values, this means merely that the system of valuations to which Jacob is committed is a resolution in something like the sense just indicated of two or more conflicting valuations or desiderata. However, precisely because the conflict is resolved, there is no genuine conflict remaining. What such strict Bayesians deny is that, if Jacob is rational, he can be committed to a system of valuations involving an unresolved conflict between rival desiderata.

Of course, Bayesians who endorse this view may concede that Jacob might not be able to tell what that resolution is. He might be able to say only that he is concerned to marry a woman who is a good cook and, at the same time, to marry a beautiful woman. He might recognize the conflict between these two descriptions of his values. Strict Bayesians would insist, however, that Jacob is committed to some resolution

of the conflict and that he would more aptly describe his values as being some resolution of the conflict between these desiderata even if he could not state precisely what that resolution is.

This response denies the presence of genuine unresolved value conflict by reducing it to a situation where Jacob has definite strict Bayesian values but is incapable of fully specifying what his commitments are. Such moves are given some legitimacy in presystematic discourse by the description of Jacob as not knowing what his goals and values are. I do not think, however, we should rely on the vagaries of ordinary language as grounds for denying the existence of genuine unresolved value conflict.

The Bayesian move would be entertainable if we construed value or utility in terms of satisfaction, desirability, preference, or in some other such utilitarian vein. Of course, even this is debatable. Ever since Plato, philosophers have noted that even our inclinations and desires can come into conflict and that it is by no means automatic that such conflicts are resolved.

Whatever the psychology of preference might be, however, utility functions need not represent only desirability or preference. They can be used to represent moral, economic, political, or various professional goals and values. I have used them to represent epistemic values appropriate to the characterization of the aims of certain kinds of inquiries. The notorious psychological facts are that men are often committed to several such systems of values and that on many occasions these may come into conflict.

Indeed, conflicts of this kind, when they occur, are often of the most serious kind in the moral life of individuals and the political life of a society. Nor are men and institutions irrational because there are conflicts between the values to which they are committed. A person who never faces such conflicts (if such a person exists) would be considered naive, simple minded, and morally insensitive (albeit ensconced in a certain sort of bliss).

Furthermore, when such conflicts arise, it is implausible to suppose that they reflect nothing more than a lack of self-knowledge on the part of the agent. If a soldier in Vietnam is called upon by his superiors to perform actions involving, in

his judgment, the gratuitous slaughter of human beings and is conflicted between his commitment as a soldier to obey the orders of superiors and his commitment to avoid the gratuitous slaughter of human beings, it would be absurd to say that the soldier's problem is to discover his own commitments. From his point of view, the challenge is not to unearth the resolution of conflict to which, unbeknownst to himself, he is already committed but to modify his commitments which are in conflict by finding a resolution toward which he can justifiably move.

The difference is not merely verbal. If the soldier were concerned only to unearth his own commitments, he might engage in the kind of self-interrogation concerning how he would make choices in various hypothetical situations which might help elicit such commitments. If he is concerned to modify his values, he will do more than that. He will attempt to justify to himself the propriety of endorsing one resolution of the conflict rather than another. The considerations relevant to such an inquiry will resemble those which arise between two parties in debate who disagree with one another over some issue of value and attempt to make a case for one point of view or another.

In cases such as this, the soldier lacks the excuse for violating Bayesian requirements that he does not know what his values are. He is, indeed, very clear what his values are and that they are in conflict. Moreover, he is also clear that he has not, as yet, found a way to resolve the conflict.

Bayesians must condemn the soldier as irrational.⁸

But this is untenable. Not only is the soldier quite rational, but he would be acting irrationally if he arbitrarily and without justification fixed on a resolution of the conflict between his value commitments just in order to save himself from the opprobrium of the Bayesians.

A rational agent should undertake inquiries to resolve conflicts in his values before making a decision—provided he is able to do so.

By the same token, it is a betrayal of reason to latch on to a resolution without justification. It is far better to acknowledge that the conflict is unresolved and to make one's decisions taking that into account.

Someone might insist, in reply, that it remains an open

question whether all modifications in value commitments are objectively arguable. Be that as it may, it seems to me that the arguability of such matters ought not to be foreclosed at the outset by a view of rational value commitment which obliges a rational agent to resolve value conflicts before choice whether such resolution is arguable or not and whether, if arguable, a justification can be offered for the resolution fixed upon.

The problem of resolving value conflicts will not be discussed in this essay. But any view which insists at the outset that the arguability of such matters is neither feasible nor relevant seems suspect. The Bayesian view of rational valuation has such an antirationalist implication and should for that reason be rejected.

8.3 Valuational Convexity

Let G be any potential commitment to a system of goals and values for possible outcomes of options feasible in a given context of choice. G is *Bayesian* if and only if it satisfies valuational consistency, closure under linear transformation, and valuational uniqueness (section 4.7).

Two Bayesian goals G_1 and G_2 *conflict* if and only if they are distinct sets of u -functions. G_3 is a *Bayesian resolution* of the conflict between Bayesian G_1 and G_2 if and only if G_3 is Bayesian and for some u_1 in G_1 and some u_2 in G_2 there is a u_3 in G_3 such that $u_3 = \alpha u_1 + (1 - \alpha)u_2$ for some $0 \leq \alpha \leq 1$.

The notion of distinct goals being in conflict should be distinguished from a single goal being in conflict. X is in goal conflict if he is "in suspense" between two or more Bayesian systems of goals and values which are in conflict.

When X is in suspense between two Bayesian systems of goals and values, he is committed to taking into account the rankings of feasible options with respect to expected utility according to both systems in assessing E -admissibility. That is to say, the conflicted system of goals and values he endorses contains u -functions of the two Bayesians goals which conflict as permissible. This is the sense in which Jacob's goals and values are in conflict in the example discussed in the previous section.

When X 's goals and values are in conflict of this sort, I assume that not only should he take into consideration all the values represented by the conflicting Bayesian goals and val-

ues but all values represented by potential Bayesian resolutions of the conflict.

If this assumption is endorsed, the set G of permissible u -functions continues to satisfy valuational consistency and closure under linear transformation; but valuational uniqueness will be violated. G will fail to be strictly Bayesian. As I have been arguing, it will be conflicted. However, G will satisfy the following condition (mentioned in chapter 4):

Valuational Convexity: If u and u' are in G then $w = \alpha u + (1 - \alpha)u'$ is also in G for every $0 \leq \alpha \leq 1$.

Unconflicted goals (that is to say, Bayesian goals) meet these three requirements. But they satisfy the stronger requirement of valuational uniqueness as well. However, conflicted goals also meet them. Hence, if rational agents are allowed to have conflicted goals and values, valuational convexity should replace valuational uniqueness as a condition on valuational rationality.

This conclusion rests on the following theses: (i) Rational agents are sometimes appropriately committed to conflicted systems of goals and values. (ii) In such conflict, u -functions which are not linear transformations of one another are permissible. (iii) All u -functions representing potential resolutions of the conflict should be permissible. (iv) A potential resolution of the conflict between values represented by distinct u -functions is a weighted average of the u -functions in conflict.

I have little to add to what has already been said in the previous section concerning (i). That (ii) is a plausible assumption about conflicted values should be apparent once the view of how an agent's goals and values are used to assess options with respect to expected utility and thereby to determine E -admissibility which has been outlined previously is adopted. To be in conflict between rival u -functions is to be incapable of coming down in favor of the one or the other for the purpose of ranking options with respect to expected utility. In such a case, ranking is done according to both utility functions, and options regarded as optimal according to at least one of them are counted as E -admissible.

Thesis (iii) also appears innocuous enough as it stands. If X has not ruled out two alternative ways of ranking options with respect to expected utility and consequences of options

with respect to his goals and values, he has not ruled out these two potential resolutions of his conflict. He should not rule out any other modes of evaluating options and consequences which qualify as potential resolutions if they exist.

What is nontrivial is the combination of thesis (iii) with a specification of the class of potential resolutions. Thesis (iv) offers such a specification.

A potential resolution of the conflict between G_1 represented by u_1 and G_2 represented by u_2 should be another Bayesian goal G_3 represented by another u -function u_3 .

It seems noncontroversial that u_3 should be uniquely determined by u_1 and u_2 up to a linear transformation. Furthermore, if we begin with linear transformations of u_1 and u_2 , the same class of linear transformations of u_3 should be determined as is determined by u_1 and u_2 . Finally, given any possible payoffs x , y , and z over which the three u -functions are defined such that $u_i(x) > u_i(y) > u_i(z)$, there is some probability value such that X would rank obtaining y for certain with obtaining x with probability p and z with probability $1 - p$ —i.e., such that $pu_i(x) + (1 - p)u_i(z) = u_i(y)$.

Armed with these conditions on potential resolutions of Bayesian goals by other Bayesian goals, one may use an argument of J. Harsanyi's designed for determining social utilities as functions of individual utilities to establish that a potential resolution must be representable by a linear transformation of some weighted average of u_1 and u_2 .⁹

Thesis (iii) becomes, in the light of this argument for thesis (iv), an injunction to take into account every weighted average of u -functions in conflict when evaluating E -admissibility. This is, of course, what valuational convexity requires.

Let G be a set of permissible u -functions over possible consequences of feasible options in some decision problem. The set G induces a quasi ordering of the possible consequences. The consequence o_{ik} is at least as highly valued as $o_{i'k'}$ according to G if and only if $u(o_{ik}) \geq u(o_{i'k'})$ for every permissible u -function in G .

This quasi ordering need not be a total ordering. Possible consequences may be noncomparable with respect to X 's system of valuations. Such noncomparability is a sure sign of value conflict.

Thus, clear-headed and rational Jacob might be able to to-

tally order his three options (in the example in the previous section) when taking into account his desire for a beautiful wife. He might do just as well when considering his interest in marrying a good cook. But since Jacob's goals and values or, if one likes, his "preferences" are in conflict, the possible consequences o_1 , o_2 , and o_3 are noncomparable. None of these outcomes is better than or equal in value to any other, all things considered.

Suppose Jacob is not conflicted but cares only to marry a beautiful woman. Strictly speaking my proposal for representing goals and values implies that Jacob is committed to a goal satisfying the strict Bayesian requirement of valuational uniqueness. However, there is no important conflict if the set of permissible u -functions consists of all those ranking o_2 over o_3 over o_1 .

The terminology of my proposal does, indeed, deviate from presystematic discourse in this respect; but the deviation seems harmless. We can distinguish between decision problems where X , if it is otherwise feasible and legitimate for him to do so, should resolve conflicts fully in order to identify a uniquely admissible option and those where it suffices to achieve only a partial resolution. In Jacob's decision problem, there is no need for further resolution of the goal conflict once he has a single total ordering of his options.

Of course, if Jacob were faced with a choice of marrying Zilpah or accepting a lottery where he has a chance p of marrying Rachel and a chance $1 - p$ of marrying Leah, the conflict due to the fact that Jacob regards as permissible every u -function ranking o_1 or o_2 over o_3 would have bite.

**8.4
Second Best
and Second
Worst**

Consider the following variants of Jacob's predicament:

Case 1: Jacob's goals and values are representable by the set G_1 of u -functions ranking o_2 over o_3 over o_1 .

Case 2: Jacob's goals and values are representable by the set G_2 of u -functions ranking o_1 over o_3 over o_2 .

Case 3: Jacob's goals and values are representable by the set G'_1 of u -functions such that $u(o_2) - u(o_3) > u(o_3) - u(o_1) > 0$.

Case 4: Jacob's goals and values are representable by the set G'_2 of u -functions such that $u(o_1) - u(o_3) > u(o_3) - u(o_2) > 0$.

Cases 1 and 3 induce the same ranking of Jacob's options with respect to expected utility. A_2 ranks over A_3 which in turn ranks over A_1 . Similarly cases 2 and 4 rank A_1 over A_2 over A_3 .

Case 5: G_3 is the *convex hull* of the u -functions in G_1 and G_2 —i.e., the set of all weighted averages of all pairs of u -functions consisting of one from G_1 and one from G_2 .

In case 5, none of the three o_i 's is comparable with any of the others. Moreover, there is some permissible ranking among the several conflicting permissible ones which ranks each of the three o_i 's best. Hence, each of the feasible options is E -admissible even though the three options are not equal in expected utility.

Case 6: G'_3 is the convex hull of the u -functions in G'_1 and G'_2 .

As in case 5, the three o_i 's (and, hence, the A_i 's) are not comparable with respect to utility. However, unlike case 5, only A_1 and A_2 are E -admissible. There is no permissible u -function which ranks A_3 on top.

In both cases 5 and 6, every u -function in G_1 (or G'_1) and in G_2 (or in G'_2) ranks marrying Zilpah second of the three alternatives. In case 5, however, there are some permissible u -functions which rank marrying Zilpah as best. Marrying Zilpah is a way of expressing a potential resolution of the conflict.

In case 6, this is not the case. That is because in both G'_1 and G'_2 marrying Zilpah is "second worst" rather than "second best" according to all u -functions. Hence, in the convex hull G'_3 , marrying Zilpah cannot express a potential resolution of the conflict. The fact that in cases 5 and 6 the quasi ordering of the options is the same fails to bring out all of the relevant differences between the two cases.

**8.5
Conflict vs.
Ignorance**

Consider the following variant on Jacob's predicament:

Case 7: Jacob is quite clear that his goals and values are representable by either G_1 or G_2 , but he cannot tell which.

This is not a case of goal conflict of the sort considered here. Jacob is not conflicted between the desire to marry a beautiful wife and to marry a good cook. He is clear that he has no such conflict. He is in doubt, however, as to which of

two rival hypotheses as to what his goals and values are is true.

My approach is no more adequate to handle Jacob's predicament in case 7 than is strict Bayesian theory. At best, we can say that marrying Zilpah is not *E*-admissible; for that option is not *E*-admissible according to either G_1 or G_2 . This stands in contrast to case 5, which might be confused with case 7, where Jacob is clear as to what his goals and values are and recognizes them to be conflicted between the desiderata of marrying a beautiful wife and marrying a good cook. In that case, all three options are clearly *E*-admissible. In case 7, Jacob knows that either marrying Rachel or marrying Leah is admissible but does not know which.

The contrast between case 5 and case 7 brings out the difference between indeterminacy in values due to lack of self-knowledge (case 7) and due to value conflict (case 5).

There are some further interpretations of indeterminacy in utility as a form of ignorance as to truth values of hypotheses which should be briefly considered.

Given any *u*-function u_i , let U_i be a corresponding truth-value-bearing hypothesis specifying for each hypothesis to which u_i assigns a utility that that utility measures the objective desirability or worth of the hypothesis.

A utility function represents a potential appraisal by an agent of truth-value-bearing hypotheses with respect to value. If that utility function (and linear transformations of it) are uniquely permissible according to *X*, *X* has a certain propositional attitude towards the hypotheses evaluated which is represented by the *u*-function.

According to one version of the view now under consideration, that propositional attitude is itself as truth-value-bearing as accepting a hypothesis in a corpus of knowledge (i.e., fully believing it) is. To count the *u*-function u_i as uniquely admissible (up to a linear transformation) is to accept the hypothesis U_i as evidence. u_i is seriously permissible according to *X* if and only if U_i is seriously possible according to *X*.

This view is clearly untenable. If U_i and $U_{i'}$ are genuinely truth-value-bearing hypotheses, *X* should be in a position to accept $U_i \vee U_{i'}$ in his corpus without accepting either disjunct. But in that event, *X*'s goals and values should count only u_i and $u_{i'}$ as permissible—counter to the convexity requirement.

Suppose someone argues that the convexity requirement should be abandoned. If *X* suspends judgment between U_i and $U_{i'}$, he should adopt a credal state over these rival hypotheses. Let his credal state be strictly Bayesian and let the degree of credence be r for U_i and $(1 - r)$ for $U_{i'}$.

It seems clear that if *X* knew that U_i is true, his utility function would be u_i and that if he knew that $U_{i'}$ is true, it would be $u_{i'}$. Hence, in his current state of knowledge, his utility function should be $ru_i + (1 - r)u_{i'}$.

Two difficulties immediately emerge:

(a) Instead of two *u*-functions being permissible (along with their linear transformations) we now have a new set of permissible utility functions. This set consists of all linear transformations of $ru_i + (1 - r)u_{i'}$. But, by hypothesis, that set corresponds to some truth-value-bearing hypothesis U^* which contradicts both U_i and $U_{i'}$. *X* cannot consistently accept U^* as evidence and suspend judgment between U_i and $U_{i'}$.

(b) By keeping r fixed and taking linear transformations of u_1 and u_2 , every weighted average of these *u*-functions becomes permissible—counter to the view under consideration.

These considerations are decisive against views which insist that valuations of hypotheses concerning the consequences of options with respect to utility themselves have truth values.

Of course, the valuations to which *X* is committed may be grounded in his knowledge and credal state. Jacob is in conflict in values because of what he knows about the three women with regard to cooking ability and physical features. This knowledge constitutes part of the warrant for the valuations he endorses. But these grounds are not truth conditions for the valuations.

Thus, the indeterminacy in valuations of hypotheses with respect to utility should not be equated with ignorance concerning the truth values of truth-value-bearing hypotheses. We have seen before that equation with ignorance of self will not do. I am now contending that ignorance about objective circumstances is no better.

The indeterminacy in such evaluations is, indeed, a sort of suspension of judgment; but it is not reducible to suspension of judgment between truth-value-bearing hypotheses. It is perhaps best to call it conflict in values.

8.6
Conflict in
Epistemic
Values

In section 2.6, I suggested that the proximate aim of any effort to expand a corpus of knowledge inductively is the acquisition of error-free information. This proximate aim was construed as involving two conflicting desiderata: avoidance of error and acquisition of new information meeting the demands of the question under investigation. The desideratum of avoidance of error can be represented by a utility function $T(g; x)$ where $T(g; t)$ is the utility afforded according to this desideratum when g is added to K and the deductive closure K_g formed when g is true and $T(g; f)$ is the utility when g is false. The utility of new information is representable by a probability measure defined over all the Boolean combinations of elements of the ultimate partition U and hypotheses which are equivalent to these given K . This information-determining probability or M -function defines the informational value of expanding to K_g from K to be $1 - M(g) = M(\sim g)$. These two utility functions represent two conflicting strictly Bayesian goals. A potential resolution of the conflict is a weighted average of these utility functions (or linear transformation thereof). In efforts at inferential expansion, I contend that we should restrict such potential resolutions to those where no error is ever ranked higher in epistemic utility to any case of avoiding error.

If we let $T(g; t) = 1$ and $T(g; f) = 0$, then

$$V(g; x) = \alpha T(g; x) + (1 - \alpha)M(\sim g),$$

where $.5 \leq \alpha \leq 1$. Here $q = (1 - \alpha)/\alpha$ represents what may be called the degree of boldness.

Thus, I have insisted all along that inferential expansion involves the resolution of conflicts. However, in my previous work, I have followed strict Bayesians in refusing to consider unresolved conflicts. In *Gambling with Truth* and other writings, I have acknowledged the existence of unresolved conflicts in epistemic values due to consideration of different lists of potential answers (different ultimate partitions) and different M -functions. I explicitly restricted the applicability of the inductive acceptance rules proposed as criteria of inferential expansion to cases where there is no such conflict.¹⁰

I was less clear, however, as to what an investigator should do in case of conflict between the demands of information occasioned by diverse questions except to suggest that he

should inquire further in the hope of resolving such conflict. In particular, it was not clear what inductive conclusions should be reached in the face of unresolved conflict.

The account of value conflict discussed in this chapter provides the basis for an answer to these questions. Epistemic value conflict can arise from three sources: conflict in the choice of a degree of boldness (or degree of caution), due to the choice of an M -function, or due to the ultimate partition used.

Let us suppose that the ultimate partition is held fixed for the moment. If the M -function is also secure and the only issue is the choice of a value for q , the proposal derivable from the considerations adduced in this chapter imply that one should take the convex hull of all epistemic utility functions derived from using the various values of q seriously considered. In section 6.2, it was shown that the P -admissible option is the one which leads to rejection of all and only elements of U which are rejected for the smallest of these q -values (or for the glb of seriously considered q -values).

If the conflict concerns the choice of M -function, one can consider the set of all epistemic utility functions obtained by considering the convex hull of M -functions initially propounded.

The criterion for evaluating expansion strategies which then emerges is one which stipulates that an element h_k of U is rejected if and only if $Q(h_k) < qM(h_k)$ for every permissible Q -function in B and every permissible M -function and for the lowest permissible value of q .

The only problem now remaining is how to handle cases where two questions are under investigation each of which has a different ultimate partition. In that event, X should adopt a new ultimate partition U^* consisting of all consistent conjunctions of one element from U_1 and another from U_2 . However, the problem is one of extending the M -functions used relative to U_1 and those used relative to U_2 to U^* .

How the extensions are to be made depends on the demands for information which occasioned the inquiries for which U_1 and U_2 generate lists of potential answers. There are, however, two kinds of extensions which, in my opinion, are legitimate.

Sometimes the inquiry for which U_1 provides potential an-

swers is such that for each $h_k \in U_1$ and for each $g_m \in U_2$ consistent given K with h_k , the informational value of h_k & g_m is no greater than the informational value of h_k and, indeed, as far as X is concerned when he focuses on that inquiry, it does not matter to him whether h_k & g_m is true or h_k & $\sim g_m$ is true as long as h_k is true; so that if h_k is true, $T(h_k \& g_m; t) = T(h_k \& g_m; f) = 1$ and $M(h_k \& g_m) = M(h_k)$. In that event, U_1 is *maximally determinate* relative to the extension to U^* .

In other cases, the demands for information are such that the T -function and M -function for U^* satisfy the requirements we have imposed previously and, in addition, are such that $M(h_k)$ is equal to what it is when U_1 is used for every h_k .

There are intermediate cases. It can be the case that relative to U^* , the h_k for some elements of U_1 are fully determinate, as in the first sort of extension, yet fail to be so for other elements of U_2 .

These intermediate cases can be represented as derivable from considering extensions of U_1 to partitions which are refinements of U_1 but coarser than U^* and where indeterminacy prevails as in the second sort of extension, and then extensions from that intermediate case to U^* which are maximally determinate.

Similar steps can be taken for U_2 . The upshot is that epistemic utility functions will be defined over U^* which, in general, will be in conflict. The convex hull can be generated and the P -admissible options determined relative to the epistemic utility functions in that convex hull.

I have outlined elsewhere how two inquiries can often be quite independent of one another, so that attending to the one without attending to the other can lead to the same conclusions when the results of both are pooled as would be the case were they considered together.¹¹

It is not my intention to elaborate on these matters any further. I mean only to point out that the account of conflict in values which has been developed can be used to good advantage in dealing with a question of some importance left unanswered in my earlier work on inferential or inductive expansion.

9

CREDAL CONVEXITY

9.1 Varieties of Ignorance

In its root meaning, ignorance is lack of knowledge. When X does not know whether h is true or false, X is ignorant concerning the truth value of h .

Such *cognitive ignorance* is part of the setting against which efforts to expand a corpus of knowledge take place. The entire point of the exercise is to acquire new error-free information and thereby to relieve some of the cognitive ignorance.

According to the account of inferential expansion developed in chapters 2, 6, and 8, which of rival expansion strategies—generated by an ultimate partition U of exclusive and exhaustive hypotheses relative to K —should be adopted depends upon X 's demands for information as represented by a (convex) set of M -functions, X 's credal state relative to K , and his index of caution q . Once such a strategy is adopted, then ignorance concerning some hypotheses will be removed. X will come to know whether g is true or false. That is to say, this will be so unless the recommended expansion strategy is not to expand at all.

Let g be equivalent given K to a disjunction of elements of U . If g is consistent, there will be some value $q(g)$ of the index q such that for $q > q(g)$, g is rejected; for $q \leq q(g)$, g goes unrejected. If g is inconsistent with K , then $q(g) = 0$. The *degree of confidence of rejection* or *degree of potential surprise* $d(g)$ is defined by $d(g) = 1 - q(g)$. The quantity $d(g)$ behaves formally like the measures of potential surprise proposed by G. L. S. Shackle.¹ Given the corpus, the demands for information, and the credal state, it measures the strength of the warrant for eliminating cognitive ignorance concerning the truth value of g by rejecting $\sim g$. In this sense, it captures an important aspect of Keynes' idea of weight of argument.²

Suppose that $d(g) = d(\sim g) = 0$. This means that no matter how bold X is prepared to be (consonant with the requirement that q be less than or equal to 1), X lacks a warrant for

removing cognitive ignorance concerning the truth value of g either by adding g to his corpus or by adding $\sim g$. Under the circumstances, X 's only recourse is to engage in further inquiry.

In such a situation, X is in a state of ignorance in a sense somewhat stronger than the sense of cognitive ignorance alone. Not only does he not know the truth value of g , but his current knowledge and his demands for information are insufficient warrant for settling the matter. I shall call this *strong cognitive ignorance*.

When X is in a state of cognitive ignorance concerning elements of a set U of hypotheses exclusive and exhaustive relative to K , X has, of course, a credal state for the elements of U .

Some Bayesians have sought to formulate a principle of inductive logic—a principle of indifference or insufficient reason—which would obligate all rational agents to assign equal Q -value to all elements of U provided that the corpus K satisfied certain conditions.

It has not, however, proven easy to formulate a principle which is at once compelling and consistent; and, for the most part, efforts to formulate such a principle of inductive logic have been abandoned. However, the idea lingers on in the conception of a state of ignorance concerning hypotheses in U which obtains if and only if the credal state for the elements of U is representable by a single Q -function assigning all elements of U equal value.

In such a state of *Bayesian ignorance*, principles of inductive logic together with the corpus K fail to obligate X to assign each element of U equal Q -value. The grounds for being in a state of Bayesian ignorance are left entirely open. Being in a state of Bayesian ignorance is simply adopting a credal state assigning each element of U equal Q -value.

I shall not quarrel with calling such states states of ignorance of some sort. Nonetheless, calling them states of ignorance is rather misleading. The credal state for U is just as definite as any other credal state for elements of U representable by a single probability function.

There is, however, another sense in which credal states can manifest ignorance of a sort which is worth considering.

In chapter 4, a contrast was drawn between those who,

while lacking any warrant for endorsing one strictly Bayesian credal state rather than another relative to what they know, choose one such state arbitrarily, and those who, also lacking the warrant, refuse to make the choice and endorse a credal state where all entertainable Q -distributions (i.e., those that have survived criticism based on considerations of inductive logic and contextual considerations) are taken into account in evaluating feasible options with respect to expected utility.

Those who pursue the latter course perform adopt numerically indeterminate credal states. They take many Q -functions to be permissible. And, in so doing, they stay in a state of suspense of a sort. Precisely for that reason, they can be said to be in a state of *credal ignorance* vis-a-vis the evaluation of Boolean combinations of elements of U with respect to credal probability.

Strict Bayesians deny credal ignorance just as they deny decision making under uncertainty. Sometimes, however, it is wise to acknowledge credal ignorance just as it is wise to acknowledge cognitive and, indeed, strong cognitive ignorance.

Credal ignorance should not be confused with cognitive ignorance. Cognitive ignorance entails suspension of judgment between rival truth-value-bearing hypotheses with respect to their truth value. Credal ignorance entails suspension of judgment between alternative systems of evaluations of hypotheses with respect to credal probability. *These modes of evaluation lack truth values*. Suspension of judgment between them is not reducible to suspension of judgment between rival truth-value-bearing hypotheses with respect to truth value.

It is very tempting to preserve a strictly Bayesian outlook by asserting the contrary. One might concede that X should sometimes be in suspense between rival Q -distributions over U . But then it might be insisted that such suspense is equivalent to suspension of judgment between rival truth-value-bearing hypotheses. There are two ways in which one might interpret this claim:

(a) One might construe the propositional attitude of allowing exactly one Q -function to be permissible as equivalent to accepting as evidence (in some suitably enriched version of L) a statement of metaphysical probability. The function Q is uniquely permissible according to X at t if and only if X accepts

as evidence in the corpus expressible in the enriched language that $M(Q)$ is the true metaphysical probability distribution where $M(Q)(h; e) = Q(h; e)$. Suspending judgment between Q_1 and Q_2 is taken to be tantamount to suspending judgment between the claim that $M(Q_1)$ is the true metaphysical probability distribution and the claim that $M(Q_2)$ is.

(b) One might construe the propositional attitude of allowing exactly one Q -function to be permissible as obtaining when X is fully aware of what his strictly Bayesian credal state is. When X is not clear as to what his credal state is, he entertains several hypotheses concerning what his strictly Bayesian state is and this accounts for the indeterminacy.

Both of these moves seek to reduce credal ignorance to cognitive ignorance, but in different ways. The first does so by translating it into ignorance concerning which of rival metaphysical probability distributions is correct. The second reduces credal ignorance to cognitive ignorance of one's own credal state.

Both approaches are untenable.

9.2 Black Boxes

I. J. Good is one of the pioneers in work on indeterminate probability judgment; his contributions are indispensable even to those who, like myself, differ as to the interpretation of what he has done.

According to Good, X 's credal state is representable by a "black box" containing a numerically precise probability function. X does not know the contents of the black box in full detail, but can make comparisons of hypotheses with respect to credal probability. Given this information and the assumption that the black-box credal state is, indeed, coherent, partial information about the contents can be gleaned and used in enabling X to behave in a coherent manner.†

† I. J. Good, "Subjective Probability as the Measure of a Non-measurable Set," in *Logic Methodology and Philosophy of Science*, edited by E. Nagel, P. Suppes, and A. Tarski, Stanford: Stanford University Press, 1962, pp. 319-329. See pp. 324-325 for an explicit statement that the black-box model proceeds as if the agent adopts a strictly Bayesian credal state unknown (Good says "unknowable") to him in all detail. On p. 327, Good introduces the idea of assigning type II distributions to hypotheses concerning the contents of the black box. Good seems to think that given a single type II distribution over alternative type I distributions (hypotheses about the contents of the black box), one can generate a new type I distribution.

Acknowledgment was made in chapter 1 that rational X would often fail to be aware of all of his commitments concerning the evaluation of hypotheses with respect to serious possibility. Yet, having acknowledged this point, attention was focused on X 's commitments and their revision. Given this concern, X was taken to be committed to knowing what his evaluations were and the issue of lack of awareness was ignored.

A similar attitude was taken towards evaluations of hypotheses concerning the consequences of feasible options with respect to goals and values. While granting that X might not be fully aware of his commitments, attention was directed to the character of the commitments.

Similarly my concern when considering X 's credal state is with his commitments. Failures to be fully aware of what these commitments are—due to emotional difficulty, limitations of memory, and computational inadequacy—are ignored.

There can be no objection, of course, to Good's directing attention to precisely those matters regarding which I have little or nothing to say. It is important, however, to appreciate the difference between examining indeterminacy in credal judgment due to lack of awareness of the contents of the black box and indeterminacy in credal judgment which is bona fide indeterminacy in the contents of the black box.

I disagree with what appears to be Good's strictly Bayesian commitment to credal uniqueness as a condition on black boxes. My reason is the one originally stated in chapter 4 and just repeated—namely, that strict Bayesian doctrine obliges X to be arbitrarily committed to a single Q -distribution even when inductive logic supplemented by contextual consideration fails to rule out many alternatives from consideration. There should be room for acknowledging genuine credal ig-

This is a confusion. Formally the type II distribution furnishes weights for the rival type I distributions and the resulting weighted average is another type I distribution. But, given the intended interpretation, the agent cannot be committed to the view that the weighted average is the correct hypothesis about the contents of the black box. Indeed, if the weighted average is different from the other type I distributions, it is not even a possibly true hypothesis about the contents of the black box from the agent's point of view. There is nothing wrong with a type II distribution provided one does not misuse it to derive such a new type I distribution.

norance without confusing it with cognitive ignorance of one's own credal states.

9.3 Metaphysical Probability

X 's credal state is a sort of propositional attitude or system of propositional attitudes adopted by X at a certain time toward a system of hypotheses. It is a system of appraisals of hypotheses with respect to credal probability.

Some propositional attitudes are intelligibly said to be true or false. Thus, X may accept h as evidence in his corpus of knowledge. We, in considering X 's corpus, can raise the question as to whether X 's belief is true or false. The belief will be true when the hypothesis h is true and false when h is false.

However, propositional attitudes do not always bear truth values. Hoping that h is true lacks truth value. Of course, a modern-day realist might toy with the idea of maintaining that X hopes that h is true if and only if X accepts the hypothesis that h is objectively hopeworthy and, hence, that hoping has a truth value after all. I hope that no one will pursue this idea and will not myself pursue it any further.

Yet, some attention should be paid to another doubtful idea. It is conceivable that someone will advance the view that when X assigns h a credal probability of r , X is accepting the hypothesis that the metaphysical probability that h is equal to r . Consequently, X 's credal judgment has a truth value.

There is already some precedent for moves of this sort. Consider the case of possibility. If the truth of h is a serious possibility according to X at time t , many will think that X accepts the hypothesis that h is objectively or metaphysically possible—i.e., that X is certain that h is true in some possible world.

My own view is that appraisals of truth-value-bearing hypotheses with respect to serious possibility and with respect to credal probability both lack truth values and that metaphysical possibility and probability ought to be assigned the same status as objective hopeworthiness.

Suppose we do say that the credal probability that h according to X is r if and only if X accepts (in his corpus in a suitably enriched language) the hypothesis that the metaphysical probability that h is equal to r .

Consider now a case where X is in suspense as to whether

the Q -value for h is r or r' . On the view being discussed, X is in suspense as to whether the metaphysical probability that h is r or r' . Because the two alternative hypotheses about metaphysical probability have truth values, X has some credal state for these alternatives. Let X assign x as the degree of credence for the hypothesis that the metaphysical probability that h is r and $1 - x$ to the alternative.

If x is different from 0 and from 1, credal coherence requires that the Q -value for h be $rx + r'(1 - x)$ which will be different from both r and r' . But this contradicts the assumption that X is in suspense between the hypothesis that the metaphysical probability is r and that the metaphysical probability is r' . To avoid contradiction, x must be either 0 or 1.

Thus, X cannot be in suspense between rival hypotheses with respect to metaphysical probability and, hence, cannot be in suspense between rival permissible Q -values for h .

Consequently, if we take appraisals of hypotheses with respect to credal probability to be assumptions about metaphysical probability which bear truth values, we are prevented altogether from allowing suspension of judgment between rival Q -functions. This is so even though no principle of logic of any sort or any other consideration might warrant our picking one assumption of metaphysical probability rather than another.

These difficulties can be avoided if we say that if X accepts the hypothesis that the metaphysical probability that h is r and his corpus contains no further relevant information (however that is to be understood), X should assign h a degree of credence equal to r . Then, however, evaluating a hypothesis with respect to credal probability will no longer be equivalent to accepting a statement of objective metaphysical probability as true and will no longer bear a truth value.

I myself consider the introduction of such metaphysical probabilities under these conditions to be gratuitous metaphysics. But no matter what one wishes to say about this, what is clear is that credal probability is not thereby seen as a truth-value-bearing mode of appraisal. One cannot see the sort of suspense between Q -functions involved in credal ignorance as reducible to suspension of judgment as to the truth of truth-value-bearing hypotheses.

**9.4
Metaphysical
Possibility**

It may, perhaps, be worth noting in passing that similar remarks apply mutatis mutandis to appraisals of truth-value-bearing hypotheses with respect to serious possibility.

Suppose someone maintains that it is seriously possible that h according to X at t if and only if X accepts (in an enriched corpus) the hypothesis that it is metaphysically possible that h .

Suppose X suspends judgment as to whether it is metaphysically possible that h or not.

X should, after all, be allowed to do so in situations where no consideration in logic, the available evidence, or context warrants favoring one alternative rather than the other.

Yet, under the conditions specified, this posture of suspense is precluded.

Consider that X must either accept h in his corpus, suspend judgment as to the truth of h , or accept $\sim h$ into his corpus. In the first and second case, the truth of h is a serious possibility according to X . Hence, X must accept "it is metaphysically possible that h ." In the third case, the truth of h is not a serious possibility according to X . Hence, X must deny "it is metaphysically possible that h ." There is no circumstance where X can be in suspense concerning the metaphysical possibility that h .

This embarrassing consequence may be avoided by weakening the link between serious and metaphysical possibility. If the truth of h is seriously possible according to X , then X accepts that it is metaphysically possible that h . The converse, however, should fail. Similarly, if X accepts that it is metaphysically impossible that h , the truth of h should not be seriously possible according to X . Once more the converse must fail.

But if this view is adopted, metaphysical possibility no longer supplies truth conditions for appraisals with respect to serious possibility. Such appraisals lack truth values. Yet evaluations of hypotheses with respect to serious possibility have an important role to play in deliberation and inquiry. Metaphysical possibility lacks such a clear role. It appears eminently expendable.

This, at any rate, is the view I shall adopt in this book.

There are, to be sure, truth-value-bearing statements of de

dicto possibility. For example, "it is possible that h according to X at t " is a truth-value-bearing report about X 's appraisals.

Sometimes we make possibility statements which are to be construed as statements concerning the consistency of hypotheses with potential corpora of knowledge. It is in this sense that I understand logical possibility—i.e., consistency with the urcorpus. To be sure, one could equate logical possibility with logical consistency—i.e., consistency with the truths of logic. But why have two terms where one will do? Unless we consider the set of logical truths to be the weakest potential corpus and, hence, consistency with the set of logical truths as a potential standard for serious possibility, there seem little benefit to the equation.

In chapter 11, I shall concede that there is, indeed, an important sense in which there are objective possibilities. But possibilities in that sense are not de dicto at all. Such possibilities are abilities and capacities such as the ability of a coin to land heads on a toss. In chapters 11 and 12, I shall also consider objective probabilities or chances and these too share some of the features of abilities. They are no more to be confused with de dicto probabilities (i.e., credal probabilities) than abilities are to be confused with serious possibilities.

**9.5
Why Convexity?**

Let h_4, h_6 , and e_H receive the interpretations they receive in case 1 of section 7.3.

Case 1: X has the following choice of options:

A_1 : If e_H is true, receive .45. If false, $-.55$.

A_2 : If e_H is true, receive $-.55$. If false, .45.

A_3 : Nothing in any case.

Let the credal state for h_4 and h_6 be described as follows: X is in suspense as to whether to consider $Q(h_4) = 0$ or $Q(h_6) = 0$ as uniquely permissible and so counts them both as permissible but nothing else.

If the credal state were so constructed, then there would be two Q -functions to use to assess expected utility. According to the first, with $Q(h_4) = 0$, the expected utilities of the three options come out as follows: .05, $-.15$, 0. When $Q(h_4) = 1$, the expectations are as follows: $-.15$, .05, 0.

The triple of numbers for each case can be viewed as having the properties of utilities assigned the hypotheses that the options are implemented. When so viewed, the suspension of judgment between the two Q -functions is manifested as a sort of conflict between two systems of values analogous to the conflicts confronting Jacob in the previous chapter.

Thus, weighted averages of the two expected-utility functions have all the earmarks of potential resolutions of the conflict; and, given the assumption that one should not preclude potential resolutions when suspending judgment between rival systems of valuations, all weighted averages of the two expected-utility functions are thus to be taken into account.

This is by no means unimportant. If only the two expected-utility functions were considered, then only A_1 and A_2 would be E -admissible. However, when the requirement is to consider all potential resolutions, then A_3 is admissible as well. If security levels are controlled by ultimate payoffs, A_3 becomes uniquely admissible.

The convexity of the credal state which requires that all weighted averages of the Q -functions assigning h_4 the values 0 and 1 guarantees the convexity of the set of expected utility functions.

As explained before, considerations such as this do not prove the convexity requirement. I have no proof. But the consideration just adduced explains how the convexity requirement fits in with the other ideas upon which this theory is being constructed.

Case 2: The options are as follows:

B_1 : If e_H , receive .55. If $\sim e_H$, $-.45$.

B_2 : If e_H , receive $-.45$. If $\sim e_H$, .55.

B_3 : Nothing in any case.

In case 2, even with convexity, only B_1 and B_2 are E -admissible. If $Q(h_4) = 0$ and $Q(e_H) = .6$, then the expected utilities are .15, $-.05$, 0. If $Q(h_4) = 1$ and $Q(e_H) = .4$, then the expected utilities are $-.05$, .15, 0.

In case 1, A_3 , the option of refusal, is, like marrying Zilpah,

a second best case in the sense that no matter which valuation scheme one uses, it is nearer the best than the worst. Hence, there are potential resolutions for which it is best. In case 2, B_3 , the option of refusal, is like marrying Zilpah when that is second worst for the conflicting modes of appraisal.

Discriminations of this sort break down if we do not mandate credal convexity.

9.6 Convexity and the Multiplication Theorem

According to credal consistency and coherence, a credal state B relative to consistent K is a nonempty set of functions each member of which is a probability measure. Let $Q_1 \in B$. The corresponding unconditional measure $Q_{1t}(h) = Q_1(h; t)$, where K entails t , has already been identified. The set B_t is the set of unconditional Q -functions obtained from the Q -functions in B in this manner.

In a similar vein, for any e consistent with K , $Q_{1e}(h) = Q_1(h; e)$ and, more generally, B_e is the set of Q_e -functions derived from Q -functions in B in this fashion.

For any Q_t -function in B_t for which $Q_t(e) > 0$, there is a unique Q_e -function in B_e such that

$$Q_e(h) = \frac{Q_t(h \& e)}{Q_t(e)}.$$

Each such matching corresponds to a distinct Q -function in B .

For those Q_t -functions for which $Q_t(e) = 0$, uniqueness no longer holds; in such cases the multiplication theorem can yield more than one associated Q_e -function. Nonetheless, each such matching still corresponds to a distinct Q -function in B .

Let W_e be the set of Q_e -functions associated with Q_t -functions in B_t via the multiplication theorem. Then credal coherence implies that $W_e = B_e$.

Credal convexity entails that for two Q_t -functions in B_t , Q_{1t} and Q_{2t} , $Q_{\alpha t} = \alpha Q_{1t} + (1 - \alpha)Q_{2t}$ is also in B_t for every $0 \leq \alpha \leq 1$.

If $Q_{1t}(e)$ and $Q_{2t}(e)$ are both positive, then the corresponding conditional Q_e -functions are given by

$$Q_{1e}(h) = \frac{Q_{1t}(h \& e)}{Q_{1t}(e)}, \quad Q_{2e}(h) = \frac{Q_{2t}(h \& e)}{Q_{2t}(e)}.$$

The multiplication theorem requires that both of these functions be in B_e .

Let V_e be the set of functions of the form

$$Q_{ae} = \alpha Q_{1e} + (1 - \alpha) Q_{2e}$$

for all $0 \leq \alpha \leq 1$.

W_e is the set of Q_e -functions of the form

$$Q_e^\alpha(h) = \frac{Q_{\alpha t}(h \& e)}{Q_{\alpha t}(e)}$$

Credal convexity entails that $V_e \subseteq B_e$.

Credal convexity for B_t and the multiplication theorem entail that $W_e \subseteq B_e$.

Furthermore, since the set of Q_t -functions consisting of all and only weighted averages of the form $Q_{\alpha t}$ of Q_{1t} and Q_{2t} form a potential unconditional credal state relative to which credal convexity entails that $V_e = B_e$, and since the multiplication theorem entails that $W_e = B_e$, we must either have $V_e = W_e$ or reject the assumption that every nonempty convex set of unconditional Q -functions can qualify as an unconditional credal state.

The matter is of some importance. It is obvious, for example, that if B_t is single-membered and $Q_t(e) > 0$, both W_e and V_e are single-membered and coincide. If it were to turn out, counter to fact, that the only circumstances under which credal coherence and credal convexity could be jointly satisfied in a nonempty credal state is when the credal state is strictly Bayesian, the effort undertaken to provide a means for representing suspension of judgment regarding such appraisals of credal probability would be in serious difficulty.

Fortunately, that is not the case. W_e must be identical with V_e for any nonempty convex B_t obeying coherence.

The proof to be offered presupposes that both $Q_{1t}(e)$ and $Q_{2t}(e)$ are positive. When one or the other is equal to 0, neither W_e nor V_e is determined by the set of weighted averages of the Q_{1t} -function and the Q_{2t} -function. We are free to guarantee their identity, therefore, without worrying about conflict with the multiplication theorem.

The proof that $W_e = V_e$ runs as follows:

For every $0 \leq \alpha \leq 1$,

$$\begin{aligned} Q_e^\alpha(h) &= \alpha Q_{1e}(h) + (1 - \alpha) Q_{2e}(h) \\ &= \alpha \frac{Q_{1t}(h \& e)}{Q_{1t}(e)} + (1 - \alpha) \frac{Q_{2t}(h \& e)}{Q_{2t}(e)}; \end{aligned}$$

$$\begin{aligned} Q_{ae}(h) &= \frac{Q_{\alpha t}(h \& e)}{Q_{\alpha t}(e)} \\ &= \frac{\alpha Q_{1t}(h \& e) + (1 - \alpha) Q_{2t}(h \& e)}{\alpha Q_{1t}(e) + (1 - \alpha) Q_{2t}(e)}. \end{aligned}$$

For any such value of α , let

$$\beta = \frac{\alpha Q_{1t}(e)}{Q_{\alpha t}(e)}.$$

It follows that

$$(1 - \beta) = \frac{(1 - \alpha) Q_{2t}(e)}{Q_{\alpha t}(e)}.$$

It also follows that $0 \leq \beta \leq 1$. Finally,

$$\begin{aligned} Q_e^\beta(h) &= \alpha \frac{Q_{1t}(e) Q_{1t}(h \& e)}{Q_{1t}(e) Q_{\alpha t}(e)} + (1 - \alpha) \frac{Q_{2t}(e) Q_{2t}(h \& e)}{Q_{2t}(e) Q_{\alpha t}(e)} \\ &= \frac{Q_{\alpha t}(h \& e)}{Q_{\alpha t}(e)} \\ &= Q_{ae}(h). \end{aligned}$$

Consequently, given any function $Q_{ae}(h)$ in V_e , there is exactly one β such that $Q_e^\beta(h)$ in W_e is identical to that function. Conversely, given any $Q_e^\beta(h)$ in W_e , one can find exactly one α such that $Q_{ae}(h)$ in V_e is identical to it. Thus the mapping between W_e and V_e is one-to-one onto, i.e., $W_e = V_e$.

In effect, what has been established in this section is that, given any convex set of Q_t -functions for which $Q_t(e)$ is positive and letting B_e be the set of Q_e -functions determined by the members of B_t via the multiplication theorem, B_e must itself be convex.

Consequently, as long as we focus on cases where $Q_t(e)$ in B_t is positive, the credal convexity condition could be formulated simply as requiring that B_t be convex. The convexity of B_e would follow automatically. Because, however, some functions in B_t may be such that $Q_t(e) = 0$, we must supply the extra stipulation that B_e be convex.

9.7 Geometrical Representations

Suppose that U is exclusive and exhaustive relative to K and each element of U is consistent with K . Also let U be finite.

If we focus attention on X 's credal state for hypotheses equivalent, given K , to elements of U and to disjunctions of elements of U , it is possible to provide a geometrical characterization of the unconditional credal state.

Let $Q_t \in B_t$ and, for each $h_i \in U$, let $x_i = Q_t(h_i)$. Since U is finite and since the Q_t -values for all Boolean combinations of elements of U are determined once the x_i 's are determined, the function Q_t can be represented by a vector consisting of $n - 1$ linearly independent components x_1, x_2, \dots, x_{n-1} . We may drop x_n because $x_n = 1 - \sum_{i=1}^{n-1} x_i$.

Because of this, the Q_t -function may be thought of as a point in an $(n - 1)$ -dimensional Euclidean space or, more narrowly, of that region in $(n - 1)$ -dimensional Euclidean space, S_t , consisting of the points whose coordinates satisfy $\sum_{i=1}^{n-1} x_i \leq 1$.

A convex set of points in the $(n - 1)$ -dimensional Euclidean space is any set of such points such that the points on the line segment joining any two of them are in the set.

Q_t is an *extreme point* of B_t if and only if there is no pair of points in B_t both distinct from Q_t such that Q_t is a weighted average of these points—i.e., such that Q_t falls on the line segment defined by these points. If B_t is a closed convex subset of S_t , it is the convex hull of its extreme points.

These ideas may be illustrated by the case where the number of elements of U is 3. In that event, the convex set S_t is the one-one right-angle triangle with right angle at the origin, as shown in figure 9.1.

In this diagram, the values of ${}^1Q_t(h_1), {}^2Q_t(h_1)$ can be read off by finding the value of the x_1 -coordinate. The same can be done for the values for h_2 by finding the value of the x_2 -coordinate. To determine the value of h_3 , one can measure the perpendicular distance to the hypotenuse and divide by $\sqrt{2}/2$.

The line segment from 1Q_t to 2Q_t represents a convex set of Q_t -functions. If the convex set includes the endpoints, the set is closed and all points in the set are boundary points. The line segment represents a potential unconditional credal state for the elements of U .

Consider all lines with a definite slope which are tangent to




Figure 9.1

the convex set. There must be at least one such line and at most two for any nonempty convex set and for any slope. Such lines are *supporting lines* for the convex set with the given slope.

When $n = 4$, we should speak of supporting planes and for $n > 4$, supporting hyperplanes.

Supporting lines parallel to the axes and parallel to the hypotenuse which is the boundary for S_t are of special interest. If there is exactly one supporting line for the convex set parallel to one of the axes (or the hypotenuse), then the unconditional credal state allows exactly one Q_t -value to be permissible for the corresponding element of U . If two such distinct supporting lines exist, they determine upper and lower Q_t -values.

9.8 Intervalism

Suppose that U contains three hypotheses as before and let the lower and upper unconditional Q -values for the three hypotheses in U be $\underline{r}(h_i)$ and $\bar{r}(h_i)$ for each h_i .

Figure 9.2 furnishes a geometric representation of the upper and lower probabilities.

The area inscribed by the bounded lines including the points on the boundary is the largest convex set with those supporting lines and it qualifies as a potential credal state relative to K . The shaded area in figure 9.3 represents, however, another




Figure 9.2

convex set with precisely the same supporting lines and, hence, with the same upper and lower probabilities.

It should be obvious from the diagram that there are infinitely many convex sets of Q_r -functions for any given set of supporting lines.

If so, a specification of upper and lower unconditional credal probabilities for each element of U fails to determine the unconditional credal state B_i uniquely.

Furthermore, supplementing these specifications with identifications of the upper and lower unconditional credal probabilities for Boolean combinations of elements of U cannot help; for the upper Q_r -value for $h_1 \vee h_2$ is one minus the lower Q_r -value for h_3 , etc.

Suppose U has $n > 3$ elements. Let U^* consist of g_1, g_2 , and g_3 , each of which is a disjunction of distinct elements of U . Specifying upper and lower Q_r -values for the g_i 's in U^* fails to determine the credal state for Boolean combinations of U^* . Let $Q(h_k; g_i)$ have a fixed value for all Q -functions in B for h_k which is a disjunct in g_i . In that event, for every permissible Q_r -function in B_i , $Q_i(h_k)$ is a maximum when $Q_i(g_i)$ is a maximum and is a minimum when $Q_i(g_i)$ is a minimum. Consequently, for such credal states over Boolean combinations of elements of U , a specification of upper and lower Q_r -values fails to determine the unconditional credal state uniquely.




Figure 9.3

B. O. Koopman introduced axioms for comparative credal probability sufficiently powerful to determine for each hypothesis in a Boolean algebra of propositions that it have an upper and lower credal probability.³ Koopman did not assume that the axioms of comparative credal probability guaranteed a total ordering of the hypotheses with respect to unconditional credal probability.

Good investigated the properties of upper and lower probabilities on the assumption that they represent the lower and upper bounds of the range of values which the unknown uniquely permissible Q -function in a strictly Bayesian black box might assign to hypotheses under consideration.⁴

C. A. B. Smith proposed a method of identifying upper and lower pigic probabilities in terms of betting behavior.⁵

H. E. Kyburg prescribed how rational agents should assign upper and lower probabilities relative to knowledge of relative frequencies.⁶

F. Schick suggested a method of extending Carnap's program to cover cases where the relative widths of predicates are not definitely determined (as they would be for Carnap, so Schick maintained, in the case of disposition predicates). His proposal involved the introduction of upper and lower probabilities.⁷

More recently, A. P. Dempster has joined the ranks of those who take lower and upper probabilities seriously.†

Some of the authors cited tend to construe indeterminate probability judgment in accordance with Good's black-box metaphor. Others tend to understand it as an expression of suspense between rival strictly Bayesian credal states.

Whatever the intention, the formal apparatuses developed could conceivably be used to represent states of indeterminate credal judgment. What are the formal and substantive similarities and differences between this mode of representing credal states and the convex-set method?

According to the convex-set method, one begins with the nonempty, convex set of Q -functions obeying the calculus of probability relative to consistent K and determines the glb and lub of permissible Q -values for each hypothesis in the algebra of propositions. These become the lower and upper probabilities for these hypotheses according to the interval-valued probability measure thus determined. This interval-valued function *envelops* the convex set in the sense just explained.

Intervalists go in the opposite direction. They begin with an interval-valued probability measure obeying some system of postulates belonging to those introduced by the authors cited above. If they are concerned about a convex set representation at all, they consider the set of all probability measures obeying the specifications of the interval-valued function. A measure obeys these specifications when the value assigned a hypothesis by the measure falls in the interval assigned that hypothesis by the interval-valued function.

The two methods are equivalent if and only if (i) permissible credal states are restricted to the largest convex sets of prob-

† A. P. Dempster, "Upper and Lower Probabilities Induced by a Multivalued Mapping," *Annals Math. Stat.*, v. 38 (1967), pp. 325-339. Dempster, to my knowledge, was the first to notice in print that upper and lower probabilities do not uniquely determine a convex set of probability measures; but Dempster credits the observation to L. J. Savage, *The Foundations of Statistics* (New York: Wiley, 1954), pp. 332-333. Dempster himself, however, endorses restrictions on credal states which guarantee that specifications of upper and lower probabilities determine the credal state to be the largest convex set compatible with the specifications. Indeed, he restricts credal states to a subset of sets constructible in this fashion. He also (mistakenly, I think) attributes to Smith the view that a credal state is representable by any convex set and, hence, is not fully characterized by upper and lower probabilities.

ability measures obeying the interval-valued functions which envelop them and (ii) interval-valued functions are restricted to functions which are envelopes of convex sets of probability measures.

Suppose condition (i) is satisfied. Suppose further than an interval-valued probability function is used to represent a credal state which is not the envelope of the largest set of probability measures obeying its specifications. In that case, (ii) is violated. Notice, however, that once the convex set is determined by the interval-valued function, the interval-valued function which does envelop the set of probability measures can be determined. If we treat the two interval-valued functions as equivalent, the violation of (ii) is nullified. Otherwise it is significant.

I contend that two credal states are similar in all respects relevant to deliberation and inquiry if and only if the same probability measures are permissible according to both credal states. This thesis is a consequence of the conception of the role of permissible Q -functions in the evaluation of E -admissibility both in practical deliberation and cognitive inquiry which I have been developing.

Given this contention, if the same set of probability measures is the largest agreeing with two distinct interval-valued functions, the two functions cannot represent distinct credal states but must be equivalent characterizations of the same credal state. On the assumption that condition (i) is satisfied, it does not then matter which of the interval-valued functions is employed from among the infinitude of such functions generally available for a given convex set. As a rule, however, it will be convenient to use the envelope function.

Some intervalists may question my contention that all relevant information about a credal state is conveyed by the set of permissible Q -functions. If so, they should state what relevant information is lost and for what purpose it is useful. Only then can we ascertain whether there are grounds for disagreement and if so what they are. For the present, I shall carry on taking my thesis for granted. The upshot is then that if condition (i) obtains, condition (ii) ought to be satisfied as well and the two methods of representing credal states are equivalent.

But should condition (i) be satisfied?

Figure 9.3 illustrates a situation where an interval-valued function envelops a convex set of probability measures without being the largest such set enveloped by the interval-valued function.

An intervalist might argue, however, that there is no relevant difference between the credal state represented by figure 9.3 and the credal state represented by figure 9.2.

Given the conception of the role of credal states in the evaluation of feasible options with respect to E -admissibility and through that with respect to admissibility, this view is mistaken. A numerical example may help settle the matter.

Let U consist of hypotheses h_1 , h_2 , and h_3 and consider B_i and B'_i for U where B_i is the convex hull of the triples of Q -values

- (a) 1/8, 3/8, 4/8
- (b) 1/8, 4/8, 3/8
- (c) 3/8, 4/8, 1/8
- (d) 4/8, 3/8, 1/8

and where B'_i is the convex hull of (a), (b), and (d). B_i is the largest convex set of unconditional probability measures enveloped by the interval-valued function assigning the following to the elements of U : (1/8, 4/8), (3/8, 4/8), (1/8, 4/8). B'_i is enveloped by the same interval-valued function.

Thus, the interval-valued representation just described cannot distinguish between the two credal states B_i and B'_i .

Perhaps, however, we have not probed far enough. Differences might emerge if we turn to the conditional credal states. That turns out not to be so. Thus, the lower and upper Q -values conditional on $h_1 \vee h_2$ for both disjuncts according to both credal states are (1/4, 4/7) and (3/7, 3/4). An examination of the two other conditional credal states reveals no difference in the lower and upper Q -values in those cases either.

Pressing on, we may still ask whether the differences between B_i and B'_i are significant for the evaluation of the admissibility of options.

Suppose X is offered a gamble on a take it or leave it basis where he wins 15 cents if h_1 is true, loses 30 cents if h_2 is true, and wins 40 cents if h_3 is true. Utility is linear in money. The

options are noncomparable with respect to strength so that the set of E -admissible and P -admissible options coincide.

When the credal state is B_i , both options are E -admissible and, hence, P -admissible. Since the security level for refusing the gamble is better than that for taking it, refusing the gamble is uniquely S -admissible and, hence, admissible.

According to B'_i , accepting the gamble is uniquely E -admissible. Hence, it is uniquely admissible.

Notice that even if we restricted attention solely to the assessment of E -admissibility, the two credal states yield different verdicts.

Thus, unless an intervalist is prepared to endorse an account of how credal states determine the admissibility of options radically different from the one proposed here, he cannot say that differences between convex sets of Q -functions enveloped by the same interval-valued function are insignificant.

One remaining avenue of defense remains open to the intervalist who endorses (i) and (ii).

Such an intervalist might argue that the principles of inductive logic restrict the set of Q -functions logically permissible relative to corpus K to the largest convex set agreeing with the specifications of a given interval-valued function. This assumption and a commitment to necessitarianism would argue in favor of following Kyburg and claiming that credal states (states of epistemological probability) be fully describable by means of interval-valued functions in accordance with condition (i).

I have already expressed my opposition to necessitarianism. But various forms of necessitarianism represent important viewpoints meriting serious critical scrutiny. At this stage of the discussion, it would be undesirable to prejudge the merits of necessitarianism one way or another. If we ruled out all credal states violating the requirements of condition (i), we would be tilting in favor of an intervalist necessitarianism. On the other hand, refusing to rule out such credal states at this stage of the discussion does not prejudice the case for doing so on the basis of considerations which may emerge later on in the discussion.

Thus, pending further reflection on the merits of necessitarianism, the representation of credal states by interval-valued functions cannot be taken to be equivalent to representation

by convex sets of probability measures. I shall continue to employ the convex-set method.

9.9 Measuring Credence

In their authoritative *Foundations of Measurement*, D. H. Krantz, R. D. Luce, P. Suppes, and A. Tversky "attempt to treat the assignment of probabilities to events as a measurement problem of the same fundamental character as the measurement of, e.g., mass or momentum."⁸ They claim that from that point of view "the debates about the meaning of probability" are about "acceptable empirical methods" of determining comparative probabilities.

These remarks seem to imply that attributing a numerical value to the mass of an object is to be understood to mean identifying the empirical methods appropriate for comparing objects with respect to whether they have more or less mass. Similarly, assigning numbers to propositions or events representing their degrees of credal or subjective probability is to be understood by identifying the empirical methods appropriate for comparing objects with respect to whether they are more or less probable.

Good used his black-box metaphor to express the idea that X 's vagueness in probability judgment is due to lack of awareness of his own strictly Bayesian credal state. Good then considered how X might secure critical control over that partial information concerning X 's strictly Bayesian credal state available to him in the form of comparisons with respect to credal probability.

Students of measurement theory seem to be concerned with two problems related to Good's:

(1) What sort of data about X 's comparisons of hypotheses with respect to credal probability would be necessary, sufficient, or necessary and sufficient to secure complete information about X 's credal state?

(2) Given the answer to (1), what conclusions of a partial sort may be obtained concerning X 's credal state from the sorts of information about X 's comparative judgments which might sometimes be practically available either to X or to another observer of X 's behavior?

As I understand Krantz et al., an answer to (1) is critical to their project. Even if it is not feasible practically to elicit sufficient data about X 's comparisons with respect to credal

probability to obtain complete information about his credal state, it should at least be feasible "in principle." Otherwise the interpretation of credal or subjective probability in terms of "acceptable empirical methods" of determining comparative probabilities will not be forthcoming.

In any case, the formal address to both questions is the same. A "relational structure," consisting of a domain of propositions and comparative probability relations on these propositions satisfying appropriate axioms, is proposed. Then the questions are asked: Is there a numerical probability measure that represents the relational structure and, if so, is it unique? If not, what is the extent of the system of probability measures qualifying as representations?

As I understand them, measurement theorists tend to resemble Good in being closet Bayesians committed to the view that a relational structure which yielded a complete characterization of X 's credal state should be representable by a unique probability measure over the domain of the relational structure.

This is revealed in the attention devoted to the problem of identifying necessary axioms for a unique representation, sufficient axioms, and necessary and sufficient axioms. More to the point, however, is the failure of such authors to devote any attention to formulating axioms characterizing indeterminate credal states.

To be sure, some measurement theorists have belatedly devoted attention to indeterminacy. Suppes has recently attempted to deal with the topic of lower and upper probabilities.⁹ He concedes that we cannot devise methods for measuring credal probabilities such that each proposition could be assigned a definite numerical probability. The situation is compared to a measurement of mass wherein the measuring instruments and procedures can yield only approximate determinations of mass, and this always with a positive error. Suppes suggests that we need a theory of approximate measurement in the case of mass and in the case of probability as well.¹⁰

Suppes proposes a theory of approximate measurement of probability that assumes all events or propositions to be comparable with respect to credal probability.¹¹ He postulates a Boolean subalgebra of the algebra of propositions generated

by n atoms all of which are equal in credal probability. Let g belong to the algebra. There must be some number k from 0 to n such that g is at least as probable as a disjunction of k atoms of the subalgebra and a disjunction of $k + 1$ elements of the subalgebra is at least as probable as g .

Sometimes g will equal one of these disjunctions in probability. In that event, the data warrant concluding that X 's assignment of a degree of credence to g is exactly k/n or $(k + 1)/n$ as the case may be. Otherwise, k/n is what Suppes considers a lower probability and $(k + 1)/n$ an upper probability.

However, Suppes explicitly states that he thinks of the upper and lower probabilities as representing the bounds of the interval within which the true credal probability for g lies.¹²

Given that this is Suppes' view of what the upper and lower probabilities are, it is unclear why he insists on considering data consisting of comparisons which yield comparability with respect to probability.

Given Suppes' prejudice in favor of strict Bayesian credal rationality, it is understandable that he should require relational structures describing credal states completely to be grounded on complete orderings of the propositions—since this requirement is necessary for the existence of a unique probability representation.

But his stated project in the aforementioned paper is to study approximate determinations, and for this there is no need to insist on comparability with respect to probability everywhere in the domain.

To be sure, the scheme proposed by Suppes will not work if full comparability is not assumed. Then we could not guarantee that there will be data to the effect that g is at least as probable as the disjunction of the k atoms of the subalgebra, nor that $k + 1$ atoms will form a disjunction at least as probable as g for some k from 0 to n .

But a scheme avoiding full comparability was proposed by Koopman a long time ago; this scheme assumes the existence of Suppes-like subalgebras for every integral value of n . In the face of the absence of full comparability, one could still determine that the strictly Bayesian credal state assigns a definite value within definite bounds to g . Koopman does not commit

himself to that interpretation of upper and lower probability, but his formalism can be bent to the purpose.¹³

S. Spielman has also developed systems of comparative probability which do not assume connected orderings. Whereas Koopman's postulation of the existence of partitions into n equiprobable alternatives suffices to guarantee the existence of at least one probability representation, Spielman supplies necessary and sufficient conditions.¹⁴

Spielman does explicitly embrace a strictly Bayesian view of the contents of the black box.¹⁵

Is it mere oversight or is it a consequence of philosophical conviction that so many students of the theory of measurement—including those who focus on question (2) and, hence, are alert to the formal issues pertaining to indeterminacy—remain faithful to strict Bayesian doctrine?

I believe that it is ideology which is decisive here. For if strict Bayesian doctrine were to be abandoned and any convex set of Q -functions relative to K taken to qualify as a credal state, measurement theorists would face a serious problem. Data concerning X 's comparisons of propositions with respect to credal probability would then be compatible with the conjecture that X 's credal state is the largest convex set of representations consistent with the ordering with respect to credal probability or with any convex subset thereof, including strictly Bayesian subsets.

That may not appear too serious. After all, if the data fail to single out a definite hypothesis as to what X 's credal state is, that circumstance is no different from what prevails even when it is assumed at the outset that whatever that state is it is strictly Bayesian.

The serious trouble is that there does not seem to be any way in which one can obtain complete information about X 's credal state from data about comparative probability when that state is one which violates strict Bayesian requirements. Consider the unconditional credal states B_i and B'_i of section 9.8, wherein h_1 , h_2 , and h_3 are all noncomparable. Any disjunction of a pair of them is no less probable than the third. The three disjunctions of pairs are noncomparable. This is true regardless of which credal state is considered. There is, in short, no way to distinguish between the credal states by appealing to comparisons of probability.

Thus if an empirical characterization of the meaning of credal probability requires a complete characterization of credal states as relational structures for comparative probability, we cannot obtain an empirical characterization.

9.10 Independence of Irrelevant Alternatives

One response to this predicament would be to reject views like the one I am advocating, which allow that agents may at least be committed to credal states which cannot be fully described in terms of comparisons with respect to credal probability on the grounds that otherwise we lack even a conception of the empirical import of what it would be like to live up to such commitments.

It is not my aim here to argue the merits of rival views concerning phenomenalism and realism, behaviorism and realism, or any other such cognate controversy. Fortunately there is little need to do so.

I have already indicated how a difference in credal states sharing the same upper and lower probability specifications can be revealed in decision making on some occasions. In the next section, I shall outline a system for ideal experimentation which could "in principle" reveal completely an agent's credal state within the framework of the theory I am proposing. The methods involve appealing to choices an agent would make in hypothetical decision problems of various sorts; but measurement theorists cannot complain about such methods since they seem prepared to use them. What cannot be obtained are complete descriptions from comparisons with respect to credal probability.

Before turning to that matter, however, it is worth drawing attention to a difference between the views of Bayesians and the approach I favor concerning decision making: this difference is important in its own right and contributes to supporting my claim that, whatever the merits of the approach towards credence and decision making I favor, it is significantly different from those of closet Bayesians such as Good, Spielman, Suppes, and other measurement theorists who speak of indeterminacy in probability.

Return to case 2 discussed in section 9.5. Suppose that, instead of facing a choice between B_1 , B_2 , and B_3 , X is compelled to choose between B_1 and B_3 . Both B_1 and B_3 would be

E -admissible and P -admissible. But B_3 is uniquely S -admissible and, hence, admissible.

However, as soon as the option B_2 is added to the set of feasible options, B_3 ceases to be admissible altogether and, what is worse yet, B_1 becomes admissible.

This circumstance stands in flagrant violation of what Luce and H. Raiffa call the condition of independence of irrelevant alternatives.¹⁶ The violation becomes more blatant if the loss from choosing B_1 when e_H is false is reduced to 44 cents so that it bears higher security level than B_2 . Then when all three options are available, B_1 becomes uniquely admissible even though when the choice was between B_1 and B_3 it was inadmissible.

Strict Bayesians must find this result anathema. If X 's credal state is strictly Bayesian and B_3 is uniquely admissible when a decision is between B_1 and B_3 , then either B_3 bears higher expected utility than B_1 or the pair of options bear equal expected utility (so that B_3 is favored because of its security). The introduction of a third option with no other relevant change in the decision problem cannot increase the expected utility of B_1 over B_3 . Indeed, given the numerical payoffs specified for the example, there is no strictly Bayesian credal state assigning e_H a definite value between .4 and .6 for which both B_1 and B_2 are optimal. And if all three options are available, B_3 cannot be optimal. Hence, if B_1 is no better than B_3 , then B_2 must be uniquely optimal. It is simply impossible for the introduction of a new option to convert an inadmissible option to an admissible one.

Furthermore, those who are convinced that an agent's choices between pairs of options reveal their rankings of these options with respect to utility or expected utility will find violation of independence of irrelevant alternatives unpalatable.

Observe, however, that if the choice is between B_1 and B_3 , both options are E -admissible. When all three options are available, B_1 and B_2 are E -admissible. The introduction of a new option does not convert an erstwhile inadmissible option with respect to expected utility into an E -admissible option. If only considerations of expected utility are considered, the independence of irrelevant alternatives is satisfied.

The truth of the matter is that B_1 and B_3 are noncomparable

with respect to expected utility. They are comparable with respect to security and, in this respect, B_3 is better. The introduction of a new option makes no difference to this. However, on the account of rational choice advocated here, admissibility is determined through the application of a lexicographically ordered sequence of criteria each of which generates its own ranking within its domain where that domain depends on the application of the criterion appearing earlier in the sequence. B_3 is better than B_1 with respect to security and noncomparable with B_1 with respect to expected utility (and strength). This is so regardless of the availability of B_2 or not. The introduction of B_2 , however, removes B_3 from the status of E -admissible, leaving B_1 and B_2 to be considered. Then the superiority of B_3 with respect to security over B_1 no longer matters. Hence, the deviation from the independence of irrelevant alternatives.

According to the approach to rational decision making I favor, X 's choice of an option from a set of feasible options does not reveal a weak preference for that option over every other option either with respect to expected utility, with respect to strength, or with respect to security. The situation is more complicated. But it is the very complexity which permits a specification of conditions under which X 's decisions, assuming he conforms to the prescriptions proposed here, could reveal his credal state.

**9.11
Supporting
Lines and the
Measurement
of Credence**

Students of the measurement of credence have, of course, utilized the work of F. P. Ramsey, B. De Finetti, and L. J. Savage to draw conclusions about X 's allegedly strict Bayesian credal state from the manner in which X makes decisions.

As we saw in section 7.3, if X is offered a gamble on e_H (the hypothesis that a given coin lands heads on a given throw) with a stake S and is asked what the lub of prices at which he would accept the bet is, the ratio P/S should be the same regardless of the magnitude of S or its sign when X 's credal state is strictly Bayesian.

On the other hand, if X 's credal state is indeterminate as in case 1 of section 7.3 the ratio will be different if S is negative than it is when S is positive. As Smith contended, when S' is equal to $-S$ for $S > 0$, then P/S will be less than or equal to P'/S' . Moreover, if the lowest permissible Q -value for e_H is \underline{r}

and the highest permissible Q -value is \bar{r} so that $\underline{r} \leq \bar{r}$, then $\underline{r} = P/S$ and $\bar{r} = P'/S'$.

Smith showed how X 's assignments of upper and lower credal probabilities could be revealed by his testimony as to how he chose in such hypothetical gambles.

To be sure, Smith's approach presupposed not only that X could adopt a numerically indeterminate credal state, but also that for gambles where the "betting quotient" is "medial" X would or should refuse the gamble.

I have avoided making strong claims about the actual behavior of all decision makers under all or even normal conditions. However, I have agreed with Smith that rational agents are often in situations where they should adopt indeterminate credal states and, moreover, that in gambles of the sort illustrated by case 1 of section 7.3, the gambles should be refused if the betting quotient is medial—i.e., that there should, indeed, be a spread in the odds.

I argued that this result could be rationalized by assuming that admissible options are S -admissible or alternatively that they were lex-admissible. In making this assumption, we have a means for extending Smith's approach to decision making in special cases exemplified by case 1 of section 7.3 to other sorts of decision problems.

The interesting implication is that, with the aid of this decision theory based on a system of criteria for admissibility lexicographically ordered in the manner proposed here, we can extend Smith's method for determining upper and lower credal probabilities to obtain a complete determination of X 's credal state.

In section 9.7, we saw that for the case where we focus on credal states for a set U of three alternatives, h_1 , h_2 , and h_3 , exclusive and exhaustive relative to K (and each consistent with it) as well as the Boolean combinations of these alternatives, for any arbitrary slope, there is at least one and at most two supporting lines to a convex set of points representing a potential credal state.

If we could determine the supporting lines for such a convex set for every possible slope, we could uniquely determine the boundary of the convex set and, hence, the convex set itself (save for whether the points on the boundary do or do not belong to the set).

As we saw, the upper and lower credal probabilities for h_1 , h_2 , and h_3 determine only three pairs of supporting lines, and this fails to determine uniquely the convex set (section 9.8).

Smith's methods provide a technique for determining the upper and lower credal probabilities of the elements of U and, hence, of their negations as well. As they stand, they fail to supply information about other supporting lines. This fact constituted the technical basis for the reservations I advanced against intervalist views of indeterminate probability judgment in section 7.8.

As I understand him, Smith did not go beyond an intervalist view. However, his techniques could have been extended to accommodate a more general approach.

Consider gambles of the following type: A stake S_1 is specified for h_1 . Another stake S_2 is specified for h_2 . X is asked to specify the lub of the prices he is willing to pay for the opportunity to take the gamble.

To appreciate the significance of the gamble, consider first the set of Q -functions for which a gamble paying $S_1 - P$ when h_1 is true, $S_2 - P$ when h_2 is true, and $-P$ when h_3 is true is "just fair"—i.e., bears 0 expectation. That set consists of points which satisfy the linear equation $x_1S_1 + x_2S_2 - P = 0$. They are representable by the points on the line segment in the right triangle of figures 9.1-9.3 of the line with slope equal to $-S_1/S_2$.

Recall now that any convex set representing a credal state has at least one and perhaps two supporting lines with the slope $-S_1/S_2$. These supporting lines determine a set of lines with slope $-S_1/S_2$ consisting of all such lines which are either supporting lines or fall between such supporting lines.

Given a gamble with payoffs determined by the stakes S_1 and S_2 and a price P which determines such a line, it must be the case that the gamble bears 0 expectation for some Q -function in the convex set representing the credal state. Hence, there are some permissible Q -functions in the credal state for which both accepting and rejecting the gamble is E -admissible and, hence, P -admissible. Hence, rejection of the gamble is uniquely S -admissible and admissible.

Consider now some line with the same slope outside the inner set. It intersects the credal state nowhere. Consequently, either it bears positive expectation for every permissible Q -

function in the credal state or negative expectation for every permissible Q -function in the credal state.

Suppose the gamble bears positive expectation. Then a gamble with stakes $-S_1$ and $-S_2$ and price $-P$ bears negative expectation for all permissible Q -functions.

Now if it should turn out that the absolute magnitude of the lub of the price at which X is willing to accept a gamble with stakes S_1 and S_2 is the same as the lub of the price when the stakes are $-S_1$ and $-S_2$ there is only one supporting line. Otherwise there are two.

Thus a procedure for identifying the supporting lines for X 's credal state with any slope has been provided. This procedure is as empirically grounded as the procedures which could be used to ascertain fair betting quotients if they exist or upper and lower credal or pignic probabilities by looking at spreads in the odds.

When $-S_1/S_2 = -1$, the supporting lines are parallel to the hypotenuse. If $-S_1/S_2 = 0$, the supporting lines are parallel to the x_1 -axis. If the ratio is infinite, they are parallel to the x_2 -axis. These three types of supporting lines are those which Smith's techniques reveal. Indeed, the method described here is equivalent to Smith's for those cases.

If the set U of exclusive and exhaustive alternatives with which we begin contains n elements, we need to determine all pairs of supporting hyperplanes for the convex set in an $(n - 1)$ -dimensional Euclidean space. These are given by constructing gambles which specify stakes for the first $n - 1$ hypotheses and ask X to indicate the lub of the prices he is willing to pay for a gamble of this kind.

Since we can "in principle" identify the supporting lines (or planes or hyperplanes) in any orientation for the convex set in the manner outlined, if we could determine all such supporting lines (planes or hyperplanes) for all orientations, we could determine X 's credal state completely. Of course, noncountably many determinations would be required so that, in practice, some method of calculation would be needed. I doubt, however, that students of the problem of measurement who follow the current orthodoxies can complain that credal states representable by convex sets which are nonempty and nonunique are inaccessible to some sort of empirical scrutiny.

Of course, this conclusion is based on the use of criteria for

rational choice quite different from the conditions which measurement theorists customarily assume are applicable. In particular, a generalization of Smith's theory has been applied which itself is based on a lexicographically ordered sequence of tests for admissibility which culminates in either a maximim or leximin criterion.

If these considerations are sound, students of the theory of measurement have no basis for objecting to the approach to credal rationality advocated here on the grounds that there would be no way in principle of ascertaining whether one was in a credal state of the sort allowed. To be sure, if agents were to satisfy the conditions for credal rationality and to make decisions in accordance with principles embedded in my prescriptions, the problem of measuring credence could not readily be formulated in the manner which measurement theorists currently favor.

I am far from prepared to claim that agents do actually satisfy the requirements of my theory with any great regularity. I simply do not know. I remain moderately confident that they could be trained to satisfy them to a moderate degree and that doing so would be better than urging them to behave like strict Bayesians. I am not deterred by the small inconvenience widespread practice along the lines I favor would cause practitioners of the current orthodox approach to measurement.

9.12 Conclusion

The case for abandoning credal uniqueness is grounded on a conviction that, wherever feasible, men should suspend judgment rather than make unwarranted and arbitrary decisions.

The substitution of convexity for uniqueness is based on a view of the use of Q -functions in evaluating options with respect to expected utility combined with the conviction that when two such systems of evaluation conflict without any warranted resolution being available, all potential resolutions should be considered permissible.

Of course, suspension of judgment incurs a cost. When credal states go indeterminate, more options are counted E -admissible and decisions become arbitrary or must be decided by considerations other than expected utility.

There is an unavoidable tension between the demand for

definiteness and the desirability of avoiding arbitrariness in credal judgment. A wise man has a just appreciation of that tension. One of my chief concerns in this book is to contribute something to an understanding of what a just appreciation of that tension entails.

10.1

Confirmational Commitments

In section 4.3, the assumption was made that rational X should always be committed to a view as to what his credal state should be relative to every potential corpus (expressible in L). That confirmational commitment was represented by a function from potential corpora to credal states. It was not assumed that every X should endorse the same confirmational commitment. That is an issue open to debate. It was also not assumed that X should remain faithful to the same confirmational commitment throughout his career as a rational agent, as the principle of *confirmational tenacity* entails. However, at any given time, X is required to be committed to some such rule.

If X is committed to $C_{x,t}$ at t , it obviously follows that $C_{x,t}(K_{x,t}) = B_{x,t}$ as the principle of *total knowledge* requires.

In section 4.3, the principle of *confirmational conditionalization* was introduced. This states that if K' is the expansion of K by adding e consistent with K and forming the deductive closure, then, for every potential confirmational commitment C , $C(K')$ is the conditionalization of $C(K)$ with respect to K and K' .

Every potential corpus is an expansion of the urcorpus UK . Hence, $C(K)$ is the conditionalization of $C(UK)$ with respect to UK and K provided that confirmational conditionalization is satisfied. $C(UK)$ is the credal state that should be endorsed in a state of complete cognitive ignorance (section 9.1) by any agent who adopts the confirmational commitment C . Probability measures of the form $P(h; e)$ in $C(UK)$ shall be called *P-functions*, to distinguish them from *Q-functions* relative to other potential corpora.

Once the set of *P-functions* in $C(UK)$ is specified, confirmational conditionalization entails that if K is obtained from UK by adding e , $C(K)$ is the set of *Q-functions* relative to K

obtained by taking each *P-function* in $C(UK)$ and constructing $Q(h; f) = P(h; f \& e)$.

Consequently, once confirmational conditionalization is adopted, confirmational commitments can be determined by identifying the set of permissible *P-functions* according to $C(UK)$.

The conditions of credal rationality discussed previously require that $C(K)$ should be a nonempty set of probability measures relative to K satisfying credal convexity, for every C and every consistent K . Confirmational conditionalization allows us to reformulate these requirements by requiring that $C(UK)$ be nonempty, that each *P-function* in $C(UK)$ be a probability measure relative to UK , and that $C(UK)$ satisfy the convexity condition.

In chapter 4, mention was made of the fact that at least two authors, H. E. Kyburg and A. P. Dempster, do not endorse confirmational conditionalization. I think they are mistaken and, hence, think it perfectly acceptable to recast the conditions of credal rationality discussed previously as conditions on $C(UK)$ and to deduce the implications for $C(K)$ from confirmational conditionalization. However, when considering the views of dissenters like Kyburg and Dempster as I shall eventually do it is important not to beg questions. For this reason, the conditions of credal coherence, consistency, and convexity have been formulated directly as requirements on credal states relative to any potential corpus and not merely on UK .

10.2

Conditional Probability

Somewhat similar considerations motivate the way in which conditional probability is treated in this book. Given confirmational conditionalization, then (for any *P-function* in $C(UK)$) $P(h; f \& e)$ defines a permissible *Q-function* relative to the corpus obtained by adding e to UK . Hence, it is tempting to think of the *P-function* itself as a representation of a strictly Bayesian rule for determining credal states relative to potential expansions of UK . If credal uniqueness is abandoned, the set of permissible *P-functions* becomes the rule.

Because conditional probability relative to potential K has customarily been considered under the assumption that confirmational conditionalization is operative, the construal of *P-functions* as furnishing instruction as to what credal states

should be adopted relative to various potential expansions of UK is quite common.

The same observation applies with some important qualification to conditional Q -functions for Q -functions in $C(K)$. If K' is a consistent expansion of K , then confirmational conditionalization implies that if $Q \in C(K)$, there is a Q' in $C(K')$ such that $Q'(h; f) = Q(h; f \& e)$ (where K' is obtained from K by adding e). Thus, conditional Q -functions relative to K can be used as rules for determining credal states relative to potential expansions of K . However, they cannot be used to determine credal states relative to all potential corpora.

Even with this qualification, however, it is easy to understand why conditional credal probability is understood as representing how, from X 's point of view, he should change his probability judgments upon the acquisition of new information.

It should be remembered, however, that conditional credal probability has an important function to play in evaluating feasible options with respect to E -admissibility. At any rate, this was the view of authors such as F. P. Ramsey, B. De Finetti, and A. Shimony. In section 5.6, a version of this view was elaborated by introducing QRCEU or MQRCEU as principles for ranking feasible options with respect to expected utility relative to specific Q -functions.

When confirmational conditionalization is endorsed, these two construals of conditional credal probability may both be endorsed without confusion. Suppose, however, confirmational conditionalization is abandoned. How should $Q(h; f \& e)$ relative to K be understood? It can be understood in its role in ranking feasible options with respect to expected utility when K is the corpus, or in its role in determining what Q' -functions should be permissible relative to K' obtained from K by adding e . These two interpretations no longer coincide. We have two notions and need two notations.

To avoid confusion, I have followed the practice of construing conditional credal probability in its functioning in evaluating expected utility. Instead of introducing another notion of conditional probability, I have directly introduced confirmational commitments as functions from potential corpora to credal states.

Instead of proceeding as I have done, one could have intro-

duced another notion of conditional probability.¹ Let $P(h; f|e)$ be understood as determining a permissible Q -function $Q(h; f)$ relative to K obtained from UK by adding e . If confirmational conditionalization obtains, then $P(h; f|e) = P(h; f \& e)$. However, if it fails, the equation might fail as well. Also, $P(h; f \& e)$ obeys the multiplication theorem. But if confirmational conditionalization is not applicable, $P(h; f|e)$ might not do so. The set of permissible functions of type $P(h; f|e)$ corresponds to $C(K)$ in my notation. If $t \in UK$, the set of permissible functions of type $P(h; f|t)$ corresponds to $C(UK)$.

I can see no objection to introducing this second notion of conditional probability and using it to characterize what I call confirmational commitments rather than employing my procedure—*provided that one does not neglect the other notion of conditional credal probability which plays an important role in ranking feasible options with respect to expected utility.*

No difficulty arises as long as confirmational conditionalization is endorsed. The two conditional probabilities are then one. I find the use of two notions and notations for conditional probability confusing when considering views which reject confirmational conditionalization—but each to his own confusion.

10.3 Harper Functions

$P(h; f|e)$, so I suggested, might represent a permissible Q -function relative to a corpus obtained by adding e to UK according to some confirmational commitment.

In a similar spirit, we might introduce $Q(h; f|e)$ for e consistent with K to represent a permissible Q' -function relative to K' obtained from K by adding e . If K itself is the deductive closure of UK and g , $Q(h; f|e) = P(h; f|e \& g)$. Hence, if confirmational conditionalization holds,

$$Q(h; f|e) = P(h; f \& e \& g) = Q(h; f \& e).$$

W. Harper has proposed a new sort of conditional Q -function which I shall represent as $Q(h; f|e)$.² When e is consistent with K , then $Q(h; f|e) = Q(h; f|e)$.

However, the Harper function is also defined for e inconsistent with K .†

† Harper ("Rational Belief Change, Popper Functions, and Counterfactuals," *Synthese*, v. 30 (1975), pt. 2, p. 1) attributes his conditional probability measures to Popper and calls them "Popper functions." This is done on the basis

In such cases $Q(h; f|e) = P(h; f|e \& w)$, where K^* is the expansion of UK obtained by adding $e \& w$ to UK and forming the deductive closure; in addition, K^* is that corpus, consistent and containing e , that most resembles K (in some unspecified sense).

One way to explicate this notion of K^* entailing a minimal revision of K is by regarding K^* as the result of first contracting K through removing $\sim e$ in a manner which minimizes loss of informational value (or implementing all contraction strategies which minimize) and then adding e to the result, thereby ending up via deductive closure with K^* .

For those who think of conditional probabilities as useful in performing the functions of confirmational commitments, Harper functions might appear to be a natural generalization from customary measures. But the generalization is quite seriously confusing.

(a) When conditional probability is understood in the sense in which it is used to rank feasible options with respect to

of appendices *iv and *v of Popper's *Logic of Scientific Discovery* (London: Hutchinson, 1959). The attribution is misleading in several respects:

(1) The idea of beginning with conditional probability as primitive is not original with Popper. Nor is the idea of doing so and permitting probability conditional on events of measure 0. I do not know who is the first to exploit both ideas, but in a paper published in Italian in 1949 and reissued in English translation as chapter 5 of his *Probability, Induction and Statistics* (New York: Wiley, 1972) De Finetti does just this.

(2) Popper, unlike De Finetti, does permit probability to be defined on impossible events. Popper does, however, acknowledge (at least indirectly) that hypotheses bearing 0 probability need not be impossible. Probability conditional on impossible events must be 1. Probability conditional on events of 0 probability need not be if the events are not impossible.

(3) Harper assumes that hypotheses bearing 0 probability, where probability represents an agent's state of belief, are inconsistent with the agent's corpus of knowledge and, hence, are, in my sense, not serious possibilities. Yet, unless they are inconsistent with the urcorpus, they remain possible for the purpose of applying Popper's theory. That is to say, if e is inconsistent with X 's corpus but consistent with UK , the unconditional probability of e will be 0 yet the conditional probability function relative to e will be a normal one with some hypotheses receiving conditional probability less than 1.

This is a special application of Popper's formalism and by no means the one Popper himself appeared to have in mind. I think it entirely appropriate to honor the inventor of the application by calling functions interpreted in the manner indicated "Harper functions."

expected utility, it makes no sense to define $Q(h; e)$ when e is inconsistent with K . Suppose, for example, that X is offered a gamble where he receives $S - P$ if $h \& e$ is true ($S > 0$), $-P$ if $\sim h \& e$ is true, and 0 if e is false. If he refuses the gamble, X receives 0. What is the maximum value of P for the fixed value of S (or the lub of such values) at which the gamble ranks higher than the alternative with respect to expected utility? The standard answer is that if e is consistent with K , then $P/S = Q(h; e)$. Hence, the conditional $Q(h; e)$ plays a significant role in evaluating the options with respect to expected utility. If, on the other hand, e is inconsistent with K so that the truth of e is not a serious possibility, the two options are ranked equally regardless of the value of P/S or of $Q(h; e)$ —assuming the latter to be defined. That is why $Q(h; e)$ goes undefined when e is inconsistent with K and conditional probability plays its role in ranking options with respect to expected utility. If we seek to preserve a notion of conditional probability which plays that role, it is desirable that that notion not be confused with Harper's conditional measures. The best way to avoid confusion is to avoid such measures if we can do without them. We can.

(b) Conditional credal probability understood as a factor in ranking options with respect to expected utility can perform that function even if the condition e has a value $Q(e) = 0$ —provided e is consistent with K . The ramifications of this point were discussed in chapter 5. There has been a tendency among students of Harper functions to overlook the difference between $Q(e) = 0$ and e being impossible (i.e., inconsistent with K).

(c) The properties one ascribes to Harper functions depend on how one thinks a shift from K containing $\sim e$ to that K^* which is a minimal revision of K consistently containing e is determined. Clarity is best served, I think, by distinguishing that matter from the question of specifying the properties of the confirmational commitment which also contribute to the characterization of Harper functions. Harper functions mask the differences between the contributions of the confirmational commitment and one's view of what constitutes a minimal revision of K .

10.4

**Irrelevant
Possibility
vs. Serious
Impossibility**

Consider the following hypothetical cases:

Case 1: X is offered a choice between the following pair of options:

A_1 : Win $S - P$ if $h \ \& \ e$ is true, win $-P$ if $\sim h \ \& \ e$ is true, and receive 0 if e is false.

A_2 : Receive 0 regardless of what is the case.

Case 2: X is offered a choice between the following pair of options:

B_1 : Win $S - P$ if $h \ \& \ e$ is true, win $-P$ if $\sim h \ \& \ e$ is true, and receive R if e is false.

B_2 : Receive 0 in any case.

In both case 1 and case 2, $h \ \& \ e$, $\sim h \ \& \ e$, and $\sim e$ are serious possibilities. The credal state for these three alternatives is identical.

Hence, in both cases, the truth of $\sim e$ is a serious possibility. Yet, as explained in section 5.6, the truth of $\sim e$ is not a relevant possibility in case 1. It is a relevant possibility in case 2.

Let case 1a be exactly like case 1 except that X adds to his corpus the information that e is true. Let case 2a be derived from case 2 in the same way.

We are going to assume that the change in knowledge does not alter X 's confirmational commitment in any way.

If this assumption is considered an unhappy one, we can consider X to have the initial corpus K and ask him about cases 1 and 2 and 1a and 2a from his current perspective when his corpus is K , his confirmational commitment is C , and his credal state is $B = C(K)$.

Relative to both 1a and 2a, the truth of $\sim e$ is no longer a serious possibility. In that case, the two cases are exactly alike. It appears entirely reasonable that A_i is admissible in 1a if and only if B_i is admissible in case 2a.

Let us now return to cases 1 and 2. $\sim e$ is a serious but irrelevant possibility in case 1 but is relevant as well as serious in case 2. The credal state may be such that option A_i is admissible in case 1 without option B_i being admissible in case 2 and vice versa.

Compare now case 1 with case 1a. The sole difference between the two cases is that what was an irrelevant but serious possibility in case 1 is no longer a serious possibility in case 1a. Initially X was guaranteed that he would neither win nor lose if $\sim e$ is true. In case 1a, such a guarantee no longer matters; for, as far as X is concerned, the truth of $\sim e$ is not a serious possibility. Does the difference between the two cases matter for admissibility?

The answer seems obvious to me. If the sole difference in the two situations is the change in the information as to the truth of e , whichever options are admissible in case 1, and only those options, should remain admissible in 1a. Since A_i is admissible in case 1a if and only if B_i is admissible in case 2a, it follows that A_i is admissible in case 1 if and only if B_i is admissible in case 2a.

Of course, what is obvious to me may not be obvious to everyone. I lack a proof that rational agents should evaluate admissibility so that the options admissible in case 1 are the same as the options admissible in case 1a. However, I confess that I am at loss to comprehend the view of anyone who differs about this matter.

To be sure, in coming to know that e is true, some other circumstance may also be altered which justifies a change in X 's confirmational commitment; but I am inviting the reader to consider situations where no such change in circumstance arises. Only in those cases do I claim that the evaluations of admissibility in cases 1 and 1a should be the same.

The idea is that to be insured against gain or loss because of the payoff structure and to be insured against gain or loss because a logical possibility is not a serious possibility are both ways of being insured against gain or loss; and it should not make any difference in deliberation which way the insurance is obtained.

If this assumption is granted, a strong case for confirmational conditionalization becomes available.

In case 1, ranking with respect to expected utility depends on the use of the Q -function conditional on e for each permissible Q -function in $B = C(K)$.

If the confirmational commitment does not change but the corpus is expanded to K' by adding e so that $C(K') = B'$,

ranking with respect to expected utility will use the unconditional Q' -functions in B .

Yet, the results of ranking should be the same—or so I am assuming. Hence, for a $Q \in B$, there should be a $Q' \in B'$ (and conversely) such that $Q(h; e) = Q'(h)$.

It would be tedious to elaborate further. It should be apparent that for potential expansions of K (which are consistent), confirmational conditionalization should be observed. In keeping with the idea that the confirmational commitment is defined for potential consistent expansions of UK , generalization of our conclusions favors imposing confirmational conditionalization throughout.

Confirmational conditionalization, however, is a restriction on potential confirmational commitments. An agent X is supposed to endorse such a commitment at time t and the rule so adopted represents what, in his view at t , should be his credal state relative to every potential corpus.

Recall, however, that in inviting the comparison between cases 1 and 1a, it was assumed that there is no relevant difference between the two cases except the change in standard for serious possibility—i.e., in the corpus of knowledge.

If there are relevant differences, the analysis might break down.

What are relevant differences? Differences in circumstance which could warrant a revision of confirmational commitment. Whether there are such relevant differences and, if so, what they are are questions which an account of the revision of confirmational commitments should attempt to answer. Some attention to these matters will be given in chapter 13.

It is important to understand now, however, that the case made for confirmational conditionalization does not warrant temporal credal conditionalization or its inverse. (See section 4.3.) The argument proves much less than some advocates of conditionalization might wish to defend.

To endorse temporal credal conditionalization and its inverse requires adoption of confirmational tenacity in addition to confirmational conditionalization—i.e., that no circumstance other than a change in corpus may warrant a revision in credal state. Such a view implies that if, in shifting from case 1 to 1a, there is no change in knowledge other than the discovery that e is true (as we have been assuming all along),

then no matter what other changes in circumstance might occur, the two cases should be treated alike. I made no such claim.

10.5 Confirmational Irrelevance

In both decision making and in scientific inquiry, X sometimes designs experiments in order to obtain data on the basis of which to make a decision or to expand his corpus via induction.

In designing such experiments to test hypotheses, investigators seek information from the data relevant to determining which (if any) of rival hypotheses to add to the body of knowledge or what practical policy to implement. Costs of reporting and storing the data argue against obtaining irrelevant information concerning the hypotheses under investigation.

Thus, if, in testing the efficacy of a drug, only the percentages of successful and unsuccessful drug therapies are relevant but data as to who was cured and who not irrelevant, it would be pointless to retain the extra information. On the other hand, if there is some doubt about relevance, it might become desirable to incur the costs of storage. In any case, judgments of relevance and irrelevance are of considerable importance to the design of experiments and the processing of information obtained from them.

The notion of relevant information in a corpus of knowledge should be distinguished from the notion of a relevant possibility discussed in section 10.4. In case 1 of section 10.4, $\sim e$ is not a relevant possibility in the context of that decision problem because the serious possibility that $\sim e$ is true may be ignored in assessing the admissibility of the available options. On the other hand, if $\sim e$ were to be added to the initial corpus, the two feasible options in case 1 would then be ranked equally even though they were ranked differently before. Moreover, the new credal state for the rival hypotheses would be modified. Adding $\sim e$ to the initial corpus would be quite relevant in these two respects.

But even when the concept of relevant information is distinguished from the notion of relevant possibility, ambiguity is possible. I am now concerned with the impact of new information on the credal state when the acquisition of that information is unaccompanied by a change in confirmational commitment. The idea is that such new information is irrele-

vant according to a given confirmational commitment when the confirmational commitment prescribes no change in the credal state through the adding of that information to the corpus.

Clearly, however, the acquisition of new information must always change the credal state in some respect—especially if the new information bore a degree of credence less than 1 prior to its acquisition.

However, in designing an experiment, we are interested in hypotheses belonging to a set U of exclusive and exhaustive hypotheses relative to an initial corpus K such that each $g_i \in U$ is consistent with K . Attention is focused on changes in the set RB of unconditional Q -distributions over the elements of U which are permissible according to the credal state $B = C(K)$ and where the changes are due to the acquisition of new information without alteration in confirmational commitment.

Let $d = e_1 \& e_2$. Let K_{e_1} and K_d be the corpora obtained by adding e_1 and d respectively to K and forming the deductive closures. Let RB_{e_1} be the set of Q -distributions over the elements of U obtained from that $C(K_{e_1})$ which is the conditionalization of $C(K)$ with respect to K and K_{e_1} ; and similarly for RB_d . Thus, if $Q \in C(K)$, then Q' is the corresponding member of $C(K_{e_1})$ and Q'' of $C(K_d)$; and $Q'(g_i) = Q(g_i; e_1)$ and $Q''(g_i) = Q(g_i; d) = Q'(g_i; e_2)$.

The sentence e_2 is *confirmationally irrelevant* to the elements of U according to the confirmational commitment C relative to K and e_1 if and only if $RB_{e_1} = RB_d$.

If e_1 is entailed by K , the relativity to e_1 may be suppressed.

If the corpus is UK , the relativity to the corpus may be suppressed.

Suppose X were to conduct an experiment and find out via routine expansion that d is true. That is to say, he expands from K to K_d while keeping his confirmational commitment constant. In reporting the results of his experiment, he might feel justified in furnishing only the information that e_1 is true if e_2 is confirmationally irrelevant. The reason is that the credal state for the rival hypotheses in U (which are the hypotheses of interest) is the same relative to K_{e_1} and relative to K_d . Insofar as X 's credal state for elements of U controls his choices or his conclusions, the information that e_2 is true can play no useful role once X finds out that e_1 is true.

In this sense, the notion of confirmational irrelevance captures an important conception of irrelevance.

It has an additional advantage. Suppose that $C(K)$ is strictly Bayesian—i.e., contains exactly one Q -function. Then e_2 is confirmationally irrelevant to the elements of U according to C relative to K and e_1 if and only if $Q(g_i; d) = Q(g_i; e_1)$ for every $g_i \in U$. This conception of confirmational irrelevance is widely adopted in some form or another by strict Bayesians. The notion I favor recognizes the strict Bayesian version as a special limiting case.

10.6 Strong Confirmational Irrelevance

There is another generalization of the strict Bayesian case which might also be considered. e_2 is *confirmationally irrelevant in the strong sense* for the elements of U according to C and relative to K and e_1 if and only if $Q(g_i; d) = Q(g_i; e_1)$ for every Q -function in B and for every $g_i \in U$.

It is obvious that strong confirmational irrelevance is sufficient for confirmational irrelevance. It is important to understand, however, that it is not necessary.

The sentence e_2 could furnish useless information as far as credal judgments about the elements of U are concerned, even though e_2 was not confirmationally irrelevant in the strong sense.

To illustrate, let U consist of g and $\sim g$, RB_{e_1} be the convex hull of $Q_{e_1}^1(g) = .6$ and $Q_{e_1}^2(g) = .4$, and RB_d be the convex hull of $Q_d^1(g) = Q_{e_1}^1(g)$ and $Q_d^2(g) = Q_{e_1}^2(g)$. Clearly $RB_d = RB_{e_1}$, so that confirmational irrelevance obtains.

Finally, let $Q_{e_1}^1(e_2) = .833$, $Q_{e_1}^2(g \& e_2) = .5$, $Q_{e_1}^2(e_2) = .75$, and $Q_{e_1}^2(g \& e_2) = .3$.

Let $Q_{\alpha e_1} = \alpha Q_{e_1}^1 + (1 - \alpha)Q_{e_1}^2$ for $0 \leq \alpha \leq 1$. Then

$$Q_{\alpha e_1}(g) = .6\alpha + .4(1 - \alpha).$$

The function $Q_{\alpha d}(g) = Q_{\alpha e_1}(g; e_2)$ is determined via the multiplication theorem to be equal to the ratio

$$\frac{Q_{\alpha e_1}(g \& e_2)}{Q_{\alpha e_1}(e_2)} = \frac{.5\alpha + .3(1 - \alpha)}{.833\alpha + .75(1 - \alpha)}.$$

Let $\alpha = .5$. Then

$$Q_{\alpha d}(g) = .506, \quad Q_{\alpha e_1}(g) = .5.$$

Thus, the condition for strong confirmational irrelevance is violated even though confirmational irrelevance obtains.

Yet, from the point of view of agent X , a gamble on the truth of g when the corpus is K_{e_1} is to be evaluated with respect to expected utility as it is to be evaluated relative to the corpus K_d . The fact that the condition for strong confirmational irrelevance is not satisfied does not matter.†

The concept of strong confirmational irrelevance is not utterly useless. It works well in the special strictly Bayesian case. It is a sufficient condition for confirmational irrelevance. Moreover, it is possible to identify a set of interesting conditions relative to which strong confirmational irrelevance becomes necessary and sufficient for confirmational irrelevance.

10.7 The Finite Case

According to Bayes' theorem (see chapter 4), every permissible Q -function in $C(K) = B$ for which $Q(e_1)$ and $Q(d)$ are positive satisfies

$$(1) \quad Q(g_i; e_1) = \frac{Q(e_1; g_i)Q(g_i)}{\sum_{j=1}^n Q(e_1; g_j)Q(g_j)}$$

$$(2) \quad Q(g_i; d) = \frac{Q(d; g_i)Q(g_i)}{\sum_{j=1}^n Q(d; g_j)Q(g_j)}$$

$$= \frac{Q(e_2; g_i \& e_1)Q(g_i; e_1)}{\sum_{j=1}^n Q(e_2; g_j \& e_1)Q(g_j; e_1)}$$

The unconditional Q -distribution over the g_i 's is called a *prior distribution* over the elements of U . Prior distributions are elements of RB .

For fixed e_1 , $Q(e_1; g_j)$ is the *likelihood* of g_j on e_1 .

$Q(g_i; e_1)$ is the *posterior* Q -value for g_i on e_1 . By confirma-

† I owe R. D. Luce a debt of gratitude for having inadvertently shown me how examples such as the one cited in the text exemplify situations where confirmational irrelevance holds but strong confirmational irrelevance fails. The observation that *confirmational* irrelevance is critical to decision making is mine. Incidentally the account of irrelevance offered here seems to me to be substantially superior to those offered by investigators who wish to take a notion of comparative probability as primitive. Their idea seems to be to require of a "representation" of a qualitative characterization of irrelevance (or "independence") that it be such that $Q(h; e) = Q(h)$. The set of such representations amount formally to a scheme for which strong confirmational irrelevance obtains. (See D. H. Krantz, R. D. Luce, P. Suppes, and A. Tversky, *Foundations of Measurement*, New York: Academic Press, 1971, v. 1, sec. 5.8.)

tional conditionalization, the posterior Q -distribution on e_1 over U is equal to a function in RB_{e_1} . Likelihoods and posterior distributions over U relative to d are defined in a similar manner.

In this discussion, U is taken to be finite and both $Q(e_1)$ and $Q(d)$ assumed to be positive.

Let us now assume that the credal state B is such that $Q(e_1; g_i) = m_i$ and $Q(d; g_i) = k_i$ for every Q -function in B . That is to say, it shall be assumed the X 's credal state is such that the likelihood function on d is the same for all permissible Q -functions and that the likelihood function on e_1 is the same for all permissible Q -functions. This is the *assumption of unique likelihoods*.

Credal coherence implies that

$$Q(d; e_1 \& g_i) = Q(e_2; e_1 \& g_i)$$

$$= k_i/m_i$$

$$= c_i.$$

The *sufficiency condition* implies that $c_i = c_i'$ for every g_i and g_i' .

Lemma I: If the sufficiency condition holds (so that the assumption of unique likelihoods obtains), then $c_i > 0$ for every g_i .

Lemma II: If the sufficiency condition holds, then, for every Q -function in B , $Q(g_i; d) = Q(g_i; e_1)$ and e_2 is confirmationally irrelevant in the strong sense for elements of U relative to K and e_1 .

Lemma III: If there is at least one Q -function in RB_{e_1} assigning positive value to every element of U and if for that function $Q(g_i; e_1) = Q(g_i; d)$ for every element g_i in U , then the sufficiency condition holds (under the assumption of unique likelihoods).

Lemma IV: If confirmational irrelevance of e_2 for elements of U obtains so that $RB_d = RB_{e_1}$, then $Q(g_i; d) = Q(g_i; e_1)$ for every extreme point of RB_{e_1} and corresponding point in RB_d and for every g_i in U .

Lemma V: Under the assumption of unique likelihoods and the condition that at least one extreme point of RB_{e_1} assigns

positive value to every element of U , lemmas III and IV imply that confirmational irrelevance entails the sufficiency condition and, hence, by lemma II, that confirmational irrelevance in the strong sense holds.

Thus, provided that $Q(e_i)$ and $Q(d)$ are positive for every Q -function, the assumption of unique likelihoods obtains, and at least one extreme point of RB_{e_i} assigns positive value to every g_i in U , strong confirmational irrelevance is not only sufficient for irrelevance but necessary as well. Needless to say, this result presupposes confirmational conditionalization as do the other results obtained in this section.

**10.8
Infinity**

Let the elements of U be represented by points with real-valued coordinates in a finite region of an n -dimensional space. No Q -function in RB , RB_{e_i} , or RB_d can assign positive value to more than a countable infinity of elements of U . (See section 5.10.) However, if every set of points with positive Lebesgue measure is assigned a positive value according to some Q -function which is an extreme point of RB_{e_i} and the assumption of unique likelihoods obtains, lemma V still applies.

Suppose U contains a countable infinity of elements. According to section 5.11, it is possible for all of them to bear 0 Q -value. Indeed, RB_{e_i} could be a convex set of such distributions and, yet, lemma V expected to apply. This will be true (given the assumption of unique likelihoods) provided at least one extreme point in RB_{e_i} is representable by a σ -finite measure assigning positive value to every element of U .

**10.9
Sufficiency
and
Irrelevance**

Suppose X knows that the chance of obtaining r heads in n tosses of coin a in some specific order is equal to $p^r(1-p)^{n-r}$ for some real value of p between 0 and 1. Any hypothesis which specifies the exact value of p is a simple-chance hypothesis asserting that coin a has a certain characteristic called a chance distribution or an objective probability distribution. The concept of objective probability or chance will be discussed at greater length in the next two chapters. In any case, we are now supposing that X considers U relative to K to consist of all hypotheses h_p specifying the exact value of p for real values between 0 and 1.

If X knows that coin a is to be tossed n times and K satisfies

conditions to be formulated in chapter 12, the principle of direct inference requires that the Q -value to be assigned the hypothesis d_r that describes the outcome of the n tosses as a specific sequence with r heads relative to the expansion K_p obtained from K by adding h_p to be equal to $p^r(1-p)^{n-r}$.

Confirmational conditionalization entails, therefore, that every Q -function in $C(K)$ should satisfy the requirement that $Q(d_r; h_p) = p^r(1-p)^{n-r}$.

If h_p is true, the chance of obtaining r heads in n tosses regardless of the order of heads and tails is equal to $\binom{n}{r}p^r(1-p)^{n-r}$. Let e_r assert that on the specific occasion of tossing coin a n times, the coin lands heads r times. Direct inference and confirmational conditionalization imply that every permissible Q -function in $C(K) = B$ satisfy the condition that

$$Q(e_r; h_p) = \binom{n}{r}p^r(1-p)^{n-r}.$$

Thus, confirmational conditionalization and direct inference secure that for fixed r , the likelihood functions for elements of U relative to e_r and d_r are unique—i.e., the same for every Q -function in B .

Furthermore, it is apparent that for each h_p , the ratio of the likelihood relative to e_r to the likelihood relative to d_r is $\binom{n}{r}$, which is the same for every value of p .

Thus, the sufficiency condition is satisfied. By lemma II, the information concerning the order of heads and tails conveyed by d_r is confirmationally irrelevant in the strong sense for U relative to K and e_r . Hence, confirmational irrelevance holds in the weak or unqualified sense as well.

Statisticians often regard the "statistic" reported by e_r to be a "contraction" of the "statistic" conveyed by d_r . Because the ratio of the likelihood of h_p given e_r to the likelihood of h_p given d_r is a constant for all values of p , the statistic reported by e_r is often called a "sufficient statistic." The sufficient statistic does not remove any relevant information from the data. As we see, this remains true according to the approach developed here.

On the other hand, suppose the contraction is not a sufficient statistic. Consider, for example, the report e_H which asserts that coin a lands heads on the first of the n tosses. By direct inference and confirmational conditionalization, every

permissible Q -function in $B = C(K)$ satisfies the requirement that $Q(e_H; h_p) = p$. Relative to d_r , it is apparent that contracting the data to e_H will lead to a violation of the sufficiency condition.

From this it follows that there are some Q -functions in $B = C(K)$ such that $Q(h_{\Delta p}; d_r)$ does not equal $Q(h_{\Delta p}; e_H)$ where $h_{\Delta p}$ asserts that the true value of p is in the interval from p to $p + \Delta p$ —namely, those where $Q(h_{\Delta p}; e_H)$ is positive. Since for some interval from some p to $p + \Delta p$ there must be at least one Q -function in the credal state for which $Q(h_{\Delta p}; e_H)$ is positive, it follows that if sufficiency breaks down, confirmational irrelevance in the strong sense also breaks down.

However, from this it does not follow that if sufficiency breaks down, e_H fails to cover all the relevant information. The extra information suppressed by reporting that e_H might be confirmationally irrelevant.

Suppose, for example, that RB (the prior credal state for elements of U) recognizes as permissible every distribution over the elements of U satisfying the requirements of the calculus of probability. Indeed, suppose we restrict attention to distributions satisfying countable additivity over the field of Lebesgue measurable sets in the interval from 0 to 1.

Under one of these circumstances, the extreme points of the convex set of Q -distributions will contain no Q -function assigning positive Q -value to every Lebesgue measurable set. And it is demonstrable that for every interval from p to $p + \Delta p$, $Q(h_{\Delta p}) = Q(h_{\Delta p}; e_H) = Q(h_{\Delta p}; e_r) = Q(h_{\Delta p}; d_r)$ for the corresponding extreme points in RB , RB_{e_H} , RB_{e_r} and RB_{d_r} ; so that these four sets are identical. Confirmational irrelevance obtains even though sufficiency breaks down.

The point being made here will reemerge later on, after the principle of direct inference has been formulated and we come to consider whether objectivist necessitarianism is tenable.

The moral here is that if the sufficiency condition is to be not only sufficient but necessary for irrelevance for elements of U , the following requirements must be met:

- (i) Confirmational conditionalization must be satisfied by the confirmational commitment.
- (ii) The likelihood functions for elements of U relative to e_1

and d must be unique (the same for all Q -functions in $B = C(K)$).

(iii) In the finite case, $Q(g_i; e_1)$ must be positive for all g_i according to at least one Q -function in B determining an extreme point of RB_{e_1} . For the other infinite cases, suitable modifications must be made as indicated in section 10.8.

Tampering with these requirements will lead to severe qualifications of the role of the sufficiency condition in determining the relevance of information conveyed by the data.

11.1
Chance and
Credal
Probability

Contrast the following two sentences:

- (i) The probability that coin *a* lands heads conditional on its being tossed at time *t* is .5 and the probability that it lands tails conditional on its being tossed at time *t* is .5.
- (ii) The probability of coin *a* landing heads on a toss is .5 and of its landing tails on a toss is also .5.

The sentence (i) expresses *X*'s appraisal of the hypotheses "coin *a* lands heads at *t*" and "coin *a* lands tails at *t*" with respect to conditional credal probability, where the condition is that coin *a* is tossed at *t*. In section 9.3, I argued that such appraisals lack truth values and, hence, cannot themselves be appraised with respect to serious possibility and credal probability.

The sentence (ii), on the other hand, is a truth-value-bearing statement. Moreover, its truth conditions do not involve reference to the cognitive states of agents but only to the state of the coin *a*. It is a statement of statistical or objective probability, rather than expression of an evaluation with respect to credal probability.

Following I. Hacking, I shall reserve the term *chance* for objective or statistical probability and shall speak of the chance of *a* landing heads on a toss rather than the probability of coin *a* landing heads on a toss.†

† I. Hacking, *Logic of Statistical Inference*, Cambridge: Cambridge University Press, 1965, pp. 10–11. In *Gambling with Truth*, I used "statistical probability" but now use Hacking's more felicitous terminology. The approach to chance I advanced in *Gambling with Truth* and shall elaborate further in this chapter and chapter 12 is heavily influenced by the ideas of R. B. Braithwaite, who defended the view that chance predicates are theoretical primitives.

In *Gambling with Truth* I compared chance predicates with disposition predicates construed as theoretical terms having a pragmatic placeholding role in investigations concerned to find explanatory systematizations. (This

The concept of chance plays an important role in scientific inquiry. In designing experiments aimed at ascertaining which of rival hypotheses belonging to some set *U* is true, an effort is made to find a "statistical model" such that each element of *U* may be equated (given the background knowledge *K*) with a statistical hypothesis specifying chances for various

idea was broached by S. Morgenbesser and myself in "Belief and Disposition," *American Phil. Quarterly* v. 1 (1964), pp. 221–232.)

By claiming that chance and disposition predicates are theoretical primitives, I mean that satisfaction conditions cannot be formulated for them in terms of descriptions of the test behavior to which they are allegedly related. As Braithwaite correctly saw, once this is conceded, the main problem in "interpreting" chance is in specifying the epistemological relation between chance and test behavior.

Hacking recognized this problem and the connection with Braithwaite's work quite independently from myself. Hacking concentrated his attention on formulating conditions relating what I call credal probabilities to objective probabilities or chances, whereas I recognized two problems in my earlier work: the problem of relating chance to credal probability and the problem of furnishing criteria for accepting hypotheses about the outcomes of testing on the basis of knowledge of chances and the inverse question of accepting (and rejecting) hypotheses about chances on the basis of knowledge of test behavior.

In *Gambling with Truth*, I was impressed chiefly with the analogies between chances and dispositions. In two interesting papers ("Possibility," *Phil. Review* v. 75 (1967), pp. 143–168 and "All Kinds of Possibility," *Phil. Review* v. 84 (1975), pp. 321–337), Hacking emphasized a parallelism between *de dicto* epistemic possibility and epistemic (i.e., credal) probability, on the one hand, and *de re* objective possibility and objective probability or chance on the other. Many of the points suggested by this parallelism are already implied by my earlier point of view, as well as by the account of knowledge as a standard for serious possibility that I developed independently of Hacking. There are also many crucial respects in which my understanding of the parallelism differs from Hacking's. Nonetheless, Hacking's papers have been of great help to me in reformulating my own view of chance and in articulating some sort of response to the currently fashionable modal realism. It would be helpful to read this chapter along with the two papers by Hacking to appreciate the points of similarity and difference.

It may be worth mentioning that I have ignored Hacking's use of grammatical tests for support in making distinctions. I do not think grammar has much relevance to the problems under discussion.

Finally, I should disclaim any commitment to Hacking's views concerning the ways in which duality with regard to possibility and with regard to probability came to be recognized. (See Hacking's *The Emergence of Probability*, Cambridge: Cambridge University Press, 1975.) I tend to be less impressed than Hacking by the role limitations imposed by "conceptual spaces" play in intellectual history, and rather suspicious of the role which the conceptual spaces of writers about such limitations play in their views on intellectual history.

“possible outcomes” of experimentation. In this manner, the problem is reduced to one of finding out which of rival statistical hypotheses is true.

It is by no means a trivial matter to construct a satisfactory statistical model; but the desirability of doing so if feasible is clear. In such cases, numerically definite likelihood functions relative to each possible outcome of experimentation are obtainable in virtue of the principle of direct inference, and the utilizability of Bayes’ theorem in deriving posterior distributions from priors in a nonarbitrary manner is at least partially achieved. (The coin example of section 10.9 illustrates this point.)

For this reason, it is important to acquire an understanding of the concept of chance sufficient to elucidate the warrant for introducing a principle of direct inference as a principle of inductive logic additional to the principle of coherence.

11.2 Ability and Serious Possibility

The concept of chance and the principle of direct inference will be examined in chapter 12. As a preliminary, some attention should be paid to the concepts of ability and compulsion; for these notions are presupposed by the concept of chance.

Sentence (ii) of section 11.1 is intended to convey (among other things) the following information.

- (a) Coin *a* is compelled to land heads or tails on a toss.
- (b) Coin *a* is incapable of landing both heads and tails on a toss.
- (c) Coin *a* is capable of landing heads on a toss and of landing tails on a toss.
- (d) Coin *a* has a .5 tendency to land heads on a toss and a .5 tendency to land tails.

These items of information are normally summed up by specifying the *kind of trial* (tossing), the *chance setup* (coin *a*), a *sample space* of possible outcomes of trials of the given kind—i.e., of kinds of responses the chance setup is capable of making on trials of the given kind (landing heads or landing tails)—that are exclusive and exhaustive, and a *probability distribution* over the sample space.

Items (a), (b), and (c) concern the specification of the chance setup, the kind of trial, and the sample space. They employ

notions of ability and compulsion. They intimate objective modality. Instead of saying that coin *a* is able to land heads on a toss, we might say that it is possible for coin *a* to land heads on a toss. In the same spirit, we can say that it is impossible for coin *a* to land both heads and tails on a toss and that coin *a* must land heads or tails (i.e., it is impossible for coin *a* to fail to land either heads or tails on a toss).

Thus, just as we have a contrast between objective chance and credal probability, it appears that we have a contrast between objective abilities and serious possibility. Moreover, the latter contrast seems in some way to lurk behind the former. Consequently, in this chapter, I shall discuss ability and compulsion (or disposition). In the next chapter I shall elaborate on the notion of tendency conveyed by (d)—i.e., on the notion of chance.

11.3 Dispositions

X knows that some iron bars attract iron filings placed nearby whereas other iron bars do not. *X* wishes to account for the difference in behavior by appealing to some trait or circumstance present in iron bars of one kind but absent in the others. That is to say, *X* seeks a predicate *D* such that the following statement is true and lawlike.

- (1) Whenever iron filings are placed near an iron bar which is a *D*, the iron filings move towards the bar.

It is not difficult to identify a predicate *D* filling the bill. Consider “is disposed to attract iron filings placed in close proximity” or “is constrained to attract iron filings placed in close proximity” or “must attract iron filings placed in close proximity.”

Of course, explaining why a particular iron bar *a* attracts iron filings placed nearby by noting that it is disposed to do so is as lame as the explanation invoking the “dormitive powers” of opium that was lampooned by Molière.

Why is the explanation a lame one? S. Morgenbesser and I suggested some time ago¹ that disposition predicates function as “placeholders” in stopgap explanations that are lame precisely because they are stopgap explanations. Whether this view implies that dispositions are not real depends on what realism is supposed to be claiming. To assert that *a* is disposed to attract iron filings placed nearby is either true or false; and

its truth value does not depend on the subjective state of any agent. In that sense, the view Morgenbesser and I advanced is congenial with realism. What we deny is that placeholding predicates, as long as they remain in that status, can be allowed to appear in lawlike statements in fully acceptable explanations. What qualifies as a fully acceptable explanation depends on the state of inquiry and the programs for explanation to which the investigators are committed. Predicates that are deficient relative to some program for explanation may, nonetheless, be used in stopgap explanations pending further inquiry; this latter will either render those predicates acceptable for purposes of explanation or will replace them with predicates that are acceptable.

The proposal to construe dispositionality as placeholding for purposes of explanation is just that—a proposal. No claim is or was made that this idea captures all aspects of presystematic usage. It does, however, seem to capture some aspects.

11.4 Disposition as Compulsion

There is one important respect in which this proposal, like most other contemporary discussions of dispositionality, is misleading. Presystematically to say that *a* is disposed to respond in manner *R* on a trial of kind *S* does not imply that *a* is compelled to respond in manner *R* on a trial of kind *S*. However, R. Carnap's view of dispositionality suggests such an implication² and, on this point, most writers have followed Carnap. According to Carnap, if *a* is disposed to respond in manner *R* on a trial of kind *S* at any time, then at that time, if *a* is subjected to a trial of kind *S*, *a* will respond in manner *R* and, in point of fact, this claim is taken not only to be true but lawlike. Yet, presystematically we do not always mean to imply that a piece of glass will break if dropped when we claim that it is disposed to break when dropped. We may imply that this is so; but we may claim only that the chances are high that it will do so.

In any case, I shall follow current practice and construe dispositionality to be invariable dispositionality or compulsion. Indeed, it is in just this sense that a predicate attributing to objects a disposition to *R* when *S*'d serves as a placeholder in stopgap explanation.

Let *N* be such a predicate and consider the statements:

- (2) Every *N R*'s when *S*'d.
- (3) Every *F R*'s when *S*'d.
- (4) All *F*'s are *N*'s.

The placeholding function of the predicate *N* is characterized by the claims:

- (5) Statement (2) is true and is a stopgap lawlike sentence.
- (6) For any predicate *F* appearing in sentences such as (3) and (4), if (3) is true and lawlike, then (4) is true and is a stopgap lawlike sentence.

In virtue of (5), it becomes clear that predicating *N* of some object is claiming that the object is compelled to respond in manner *R* whenever subjected to a trial of kind *S*.

I shall call any placeholding predicate of this sort a disposition predicate or a predicate of compulsion or a predicate of necessitation or, finally, an *N*-predicate.

Suppose an *N*-predicate is denied to be true of some object *a*. The object *a* is not compelled to *R* when *S*'d. In that sense, *a* is able to fail to *R* when *S*'d. It is possible for *a* not to *R* when *S*'d. I shall call the negation of an *N*-predicate a predicate of ability or possibility or a *P*-predicate.

11.5 Explicit *N*-Predicates

Some disposition predicates or predicates of compulsion wear their status on their sleeves. Others wear their status covertly; and others fail to do so in any obvious manner at all. Consider the difference between "is compelled to dissolve when immersed in water" and "is water soluble" or between "is disposed to attract iron filings placed nearby" and "is magnetic." In both cases, the first term in the pair is explicitly dispositional whereas the second term is less so.

The difference carries some small importance. If an *N*-predicate is explicitly so in the sense that it clearly indicates the kind of trial and outcome with respect to which it is dispositional as well as indicating on its face that it is dispositional, the claim that (2) is true is incorrigible in an important sense as is the claim that if (3) is true so is (4). Moreover, the stopgap lawlike status of (2) and (4) cannot be eliminated. This is not so if an *N*-predicate is not explicitly so.

Thus, it is entirely envisageable that the "reduction sentences" linking "is magnetic" with test behavior may be revised in important ways and that the predicate "is magnetic" might become so acceptably integrated into scientific theory that it loses its placeholder status. This is not so for "is compelled to attract iron filings placed nearby."

To understand the difference, let "is an $N(R/S)$ " be an explicit compulsion predicate to be paraphrased as "is compelled to respond in manner R when S 'd." Imagine that a given text contains the following sentence:

(7) a is an N .*

At the bottom of the page, there appears a footnote:

*(5) is true and so is (6).

We could, however, contemplate modifying the footnote by an abbreviation understood by everyone:

*(R/S).

Finally, we could eliminate the footnote by converting it to an index on the predicate in the sentence (7) so that we have the following:

(7') a is an $N(R/S)$.

The point of this exercise is to suggest that the prima facie structure in the explicit compulsion predicate "is compelled to respond in manner R when S 'd" is to supply background information concerning true lawlike statements in which the N -predicate appears. Anyone who is prepared to use the explicit N -predicate is thereby committed to assuming the alleged background information as part of his corpus of knowledge. In this sense, it may be regarded as part of the incorrigible urcorpus UK for the language in which the predicate "is an $N(R/S)$ " appears.

In the case of predicates like "is magnetic" such background information is not explicitly supplied in using the predicate and that information can be modified while the use of the predicate is retained.

Similar remarks apply mutatis mutandis to predicates of ability. There are explicit predicates of ability of the form "is a $P(R/S)$ "—i.e., "is able to R when S 'd."

11.6 Ability Adverbially Modified

Suppose "Everything R 's when S 'd" is true and lawlike.

In such cases, there is no point in introducing a placeholding predicate "is an $N(R/S)$ " for purposes of explanation. If we explain why something R 'd by citing the law and the fact that the object was S 'd, there is no point in adding the further gloss that the thing was compelled to R when S 'd.

Yet, when we have a linguistic device to generate new predicates out of old which has a function in some contexts, there is a temptation to use it even in contexts where it has no function. As long as we understand what is happening, there is little harm in it. We should not be deceived into thinking we have added anything to our knowledge when we claim that everything is compelled to R when S 'd.

Attributions of abilities to agents and things are often qualified in various ways. Thus, Levi is physically incapable of running a four-minute mile. It is physically possible for Levi to spend a fortune at the roulette table at a Las Vegas Casino but it is not economically possible for him to do so.

The claim that it is physically impossible for Levi to run a four-minute mile is equivalent to the contention that he is physically compelled to fail to run a four-minute mile in an effort on his part to do so. The adverb "physically" is an explicit indication that the description of Levi's traits and circumstances which should be sought as a replacement for the placeholding N -predicate should be constituted out of terms of physical theory (however that is to be understood). Much the same can be said for "biologically," "psychologically," "economically," et al. To say that it is technologically possible (or impossible) is not to allude to replacement by terms appropriate to some specific domain, theory, or science but to promise replacement by descriptions of the agent's knowledge or ignorance and the resources available to him.

Thus, to assert that it is physically impossible for Levi to run a four-minute mile is not to claim that it is inconsistent with the true laws of physics that Levi runs a four-minute mile. To my knowledge, there is no such inconsistency. However, the true statement that Levi is incapable of running a four-minute mile together with the appropriate stopgap law of type (2) entails that Levi does not run a four-minute mile. The attribution of physical impossibility here is an attribution of physical incapacity.

I have no objection to introducing the notion of physical impossibility as inconsistency with the true laws of physics. Observe, however, that so-construed physical impossibility is not predicated of things and agents but of statements. One might claim that it is physically possible *that* Levi run a four-minute mile at time *t* and place *p* because it is consistent with the true laws of physics that this be so. Yet, we can deny that Levi is capable of running a four-minute mile at time *t* and place *p*. That is to say, we can deny that it is physically possible *for* Levi to perform such a feat.

Once this is understood, we may coherently speak of a contrast between physical possibility that and physical possibility for, biological possibility that and biological possibility for, et al. But whereas there is good sense in claiming that it is technologically possible for an agent to do so and so, it makes no sense to say that it is technologically possible that the agent do so and so—unless one insists that there are laws of technology.

Consider now logical possibility. It makes sense to speak of logical possibility that *h*; for that amounts to no more than logical consistency. What about logical possibility for an agent or object to do something?

Some authors appear to think that objects may be logically compelled or constrained to certain kinds of behavior just as they may be compelled, disposed, or constrained physically or biologically to other kinds of behavior.† I find such views incomprehensible.

If we are to follow the analogy with physical possibility or physical compulsion, to say that *a* is logically compelled to *R* when *S*'d is to say that *a* is an *N(R/S)* and to promise replacement of the *N*-predicate by a description in purely logical terms. I cannot make very much sense out of this. Perhaps, what is intended is that the variant of (2) becomes a logical truth when the placeholder is replaced by the new description. But if "everything *R*'s whenever it is *S*'d" is not a logical truth, then (2) cannot be a logical truth unless the substitution for "is an *N(R/S)*" is inconsistent.

† Hacking, "All Kinds of Possibility," *Phil. Review*, v. 84 (1975), pp. 332-334. In all fairness to Hacking, he recognized the bizarre character of his conclusion and put it forward only because of the encouragement he found in Kripke's work.

We can obtain a sort of logical compulsion if, in those cases where "everything *R*'s whenever it is *S*'d" is a logical truth, we introduce the predicate "is an *N(R/S)*" with the background stipulation in lieu of (5) and (6) as follows:

(8) "Everything is an *N(R/S)*" is true and lawlike.

No harm arises from using this device provided its operation is well understood; but there is no great benefit in it either. If one wishes, it can be said to be a predicate of logical compulsion and its negation a predicate of logical possibility—i.e., logical ability.

N-predicates are normally introduced as placeholders for conditions under which processes or behaviors of some kinds occur once other processes have occurred. There is no need, therefore, for *N*-predicates of *N*-predicates. We do not need "is compelled to *N(R/S)* on a trial of kind *T*" or "is an *N(N(R/S)/T)*."

On the other hand, we do acknowledge that some objects are magnetizable in the sense that they are compelled to become magnetic when subjected to a trial of a given kind. We can have "is compelled to become an *N(R/S)* on a trial of kind *T*." Such *N*-predicates do not illustrate iteration of the *N*-operator.

Carnap's account of disposition predicates implies that conducting a trial of kind *S* and obtaining a result of kind *R* is sufficient warrant for attributing the disposition to *R* on a trial of kind *S*. This is a mistake. Coin *a* is tossed once and lands heads. *a* is capable of landing heads on a toss; but it is not compelled to land heads on a toss nor is it invariably disposed to land heads on a toss.³

11.7 Disposition as Primitive

Some will complain that my treatment of ability and compulsion does not qualify as analysis because it fails to specify necessary and sufficient satisfaction conditions for such predicates in terms of test behavior.

No such semantics has been offered; but no such semantics (and, hence, no such analysis) is needed. Once the commentary for a given *N*-predicate is specified, *X* has instructions as to how to evaluate hypotheses with respect to serious possibility concerning the test behavior of objects of which the *N*-predicate is known to be true.

Furthermore, X 's corpus, confirmational commitment, and the commentary will provide a basis for X to reach conclusions as to whether statements of compulsion and ability should or should not be added to his corpus. Thus, if a is known to have responded without fail in manner R on a great many trials of kind S , X might have a credal state which assigns credal probability to the hypothesis that a is able to fail to R on a trial of kind S sufficiently low to warrant—given his demands for information—rejecting that hypothesis.

Of course, it may be desirable to find necessary and sufficient satisfaction conditions for N -predicates. Indeed, their function as placeholders for purposes of explanation suggests that it is eminently desirable to find alternative equivalent descriptions which are explanatorily adequate; but this is tantamount to looking for necessary and sufficient satisfaction conditions of a certain sort.

The required conditions might turn out to be specifications of microstructures of the objects having the disposition rather than descriptions of test behavior (although microstructural descriptions need not be obtained in order to integrate an N -predicate into a theory.)

Thus, I am not denying that predicates of inability and ability often need a semantics; but the semantics wanted is not to be obtained by invoking a possible-worlds semantics. Nor is to be sought by tampering with "logical form" so as to cosy up to convention T without infinitely many primitives. Rather it is to be sought through the conduct of special inquiries into the subject matters for which the predicates of ability and inability are used.

We improve our understanding of predicates such as "is compelled to attract iron filings placed nearby" by studying magnetic theory, and not by studying possible worlds or any other armchair semantics.

**11.8
Knowledge of
Ability and
Appraisal of
Serious
Possibility**

Suppose X confronts an urn b containing a large number of coins. X knows that some of the coins are two-headed and some are two-tailed and none have heads on one side and tails on the other. The urn b and its contents are the "setup." Let a trial of kind S be drawing a coin from the urn while blindfolded and after the contents have been thoroughly mixed and

tossing the coin twice. Let R describe the outcome of the coin landing heads on the second toss in a trial of kind S .

X knows that b is able to respond in manner R on a trial of kind S .

Let a trial of kind T be one where a trial of kind S is performed and the coin lands tails on the first toss.

X knows that b is incapable of responding in manner R on a trial of kind T .

Suppose X knows that at t a trial of kind S has occurred.

Is the hypothesis "a result of kind R occurs at t " a serious possibility according to X ?

This question broaches the issue of how knowledge of objective possibility and impossibility—i.e., ability and compulsion—controls or should control appraisals of hypotheses with respect to serious possibility.

It is useful to consider this question as a preliminary to the question of direct inference. This last question concerns the conditions under which judgments of credal probability are grounded on knowledge of objective chance. The question now under consideration concerns the conditions under which judgments of serious possibility are grounded on knowledge of objective compulsion and ability.

At the outset, we should remember the considerations aduced in section 9.4 against construing appraisals of hypotheses with respect to serious possibility as bearing truth values. It is not the case that h is seriously possible according to X at t if and only if X accepts as evidence that it is metaphysically possible that h .

On the other hand, under certain circumstances, X 's knowledge that it is possible for a result of kind R to occur on a trial of kind S and that a trial of kind S has occurred warrants his evaluating the truth of the hypothesis that a result of kind R occurs as a serious possibility. In this sense, knowledge of objective possibility can justify an evaluation of a hypothesis with respect to serious possibility.

Observe, however, that the knowledge of objective possibility involved is knowledge that it is objectively possible for a setup b to respond in a certain manner R on a trial of kind S and not knowledge that it is objectively (i.e., metaphysically) possible that b respond in a certain manner R at time t . Pos-

sibility for is ability. The knowledge is that b is capable of responding in manner R on a trial of kind S .

Furthermore, this knowledge is not equivalent to appraisal of the hypothesis that b responds in manner R at t as seriously possible.

Indeed, the link between the knowledge of objective possibility or ability and the appraisal with respect to serious possibility is quite fragile. It depends on X not knowing too much.

Suppose, for example, that X also knows that the trial of kind S under consideration is also a trial of kind T . Then even though he continues to know that it is possible for b to respond in manner R on a trial of kind S , he no longer evaluates " b responds in manner R at t " as seriously possible. He considers it as not seriously possible in virtue of his new knowledge and the knowledge he already has that b is compelled not to R on a trial of kind T .

Thus, the grounding of appraisals of hypotheses about outcomes of trials with respect to serious possibility on knowledge of abilities or objective possibility depends on X 's other knowledge. The additional knowledge at his disposal can prevent evaluating such a hypothesis as a serious possibility even in the face of knowledge of ability or objective possibility.

The point is familiar enough. However, the analogy to direct inference is not always appreciated. In grounding judgments of credal probability on knowledge of chances or objective probabilities, it is crucial that care be taken that X 's other knowledge not prevent a given judgment from going through.

Thus, if X knows that the chance of heads on a toss of coin a is .5, everything else being equal, knowledge that coin a is tossed should warrant his assigning credence of .5 to the hypothesis that the coin lands heads. But "everything else being equal" involves X not knowing too much or too little.

11.9 Subjunctive Conditionals

Disposition statements and statements of ability support subjunctive conditionals. What does "support" mean here?

Counterfactual conditionals of the type "if h were true, g would be true" are expressions of modal appraisal just as statements such as "it is possibly true that h " are. Since such unconditional modal appraisals lack truth values, so do counterfactual modal appraisals.

Counterfactual modal appraisals are made relative to a corpus where the hypothesis h in the antecedent is known to be false. $\sim h \in K$.

Such appraisals are evaluations of hypotheses occurring in the consequents of the conditionals with respect to possibility relative to corpora that are transformations of K of certain kinds. The corpus K_2 is obtained from K by first contracting to K_1 through removing $\sim h$ in a manner which minimizes loss of informational value. (See chapter 3.) Then K_2 is the deductive closure of K_1 and h .

Open subjunctive conditional appraisals are made when neither h nor $\sim h$ is in K . The idea is then to obtain K_2 by adding h to K and forming the deductive closure.

In both the counterfactual and open conditional appraisal, the consequent is that g would be true just in case g is entailed by K_2 and is that g could be true if g is consistent with K_2 .

It is important to understand that neither form of subjunctive appraisal has a truth value. Consequently, subjunctive conditionals are not to be evaluated themselves with respect to serious possibility or credal probability. The recent controversy concerning whether the probability of a conditional is a conditional probability or not is one of the more egregious examples of a nonissue emerging from modal metaphysics.⁴

To illustrate the way this view applies to the connection between knowledge of ability and inability, suppose X knows that the following are true:

(9) a is an $N(R/S)$.

(10) A trial of kind S is not performed on a at t .

In addition, X knows the background information contained in the abbreviated gloss " (R/S) "—in particular, that (2) is true and that the following is true:

(11) An event of kind R does not occur at a at t .

X now considers the counterfactual appraisal expressed by "if a were S 'd at t , it would have R 'd." The contracted corpus K_1 involves removing (10) in a manner minimizing loss of information. It is a condition of the problem that (9) remain. Hence, either (11) or (2) must go. But (2) is lawlike and the loss of information incurred relatively substantial. (11) is the obvious candidate for removal by this criterion. Expansion to

K_2 by adding the contradictory (10) brings in the denial of (11), and the evaluation expressed in the counterfactual is supported.

Among the many virtues of this approach to subjunctive conditionals is that the notion of similarity between possible worlds is replaced by evaluations of contractions of K through removing $\sim h$, which minimizes the loss of information. Appraisals of potential contraction strategies with respect to information losses become the key to handling the question of cotenability. The role of informational value is going to be central to an account of the revision of knowledge anyhow. It is a pleasant fact that it can contribute to our understanding of cotenability in a manner which eliminates the gratuitous mysteries of modal metaphysics.

An obvious implication of the approach I favor is that one should abandon efforts to provide satisfaction conditions for predicates of compulsion and ability in terms of counterfactual or other subjunctive conditionals. One may go a small step further. Any effort to explicate disposition predicates in terms of the counterfactuals they support appears to put the cart before the horse.

The way to explicate predicates of compulsion and ability is to identify the postulates appearing in the commentary associated with such predicates. That commentary and the pragmatic role provide all one can expect in the way of "conceptual analysis" of such predicates. The point is a small one. It looms a little larger when we turn to a discussion of chance.

11.10 Sample Spaces

Return to sentence (ii) of section 11.1. In section 11.2, I claimed that sentence (ii) contained the four items of information (a), (b), (c), and (d). The first three conditions specify the setup a of which the chance predicate is predicated, the kind of trial, and the sample space. In specifying the sample space, these conditions assert that the chance setup is compelled to respond in certain ways on the kind of trial and is able to respond in other ways.

I suggest that we introduce the technical device of a *CD-predicate* of the type "is a $CD(\mathcal{F}(\Omega)/S)$." This sort of predicate is a complex placeholding predicate. To explain it, consider a family of predicates of the type R_α where α is a standard designator for a subset of a set Ω of n -tuples of real numbers

belonging to a Boolean or σ -algebra $\mathcal{F}(\Omega)$ containing all unit sets and Ω .

The expression " $(\mathcal{F}(\Omega)/S)$ " is to be understood as an abbreviation for the commentary:

- (a) All CD 's are $N(R_\alpha/S)$'s.
- (b) Given any α and β such that $\alpha \cap \beta = \emptyset$ and $\alpha \in \mathcal{F}(\Omega)$ and $\beta \in \mathcal{F}(\Omega)$, then "All CD 's are $N(\sim(R_\alpha \& R_\beta)/S)$'s" is true.
- (c) If α is a unit subset of $\mathcal{F}(\Omega)$, then "all CD 's are $P(R_\alpha/S)$'s" is true.

The predicate S and the various predicates of the type R_α are predicates true or false of events.† Notice that if a is a $CD(\mathcal{F}(\Omega)/S)$, the marginal commentary does not guarantee that a is a $CD(\mathcal{F}(\Omega)/S \& T)$. This inference is prevented by condition (c). In the case of the tossed coin, (c) asserts that it is possible for the coin to land heads and for the coin to land tails on a toss. But suppose a trial which is both S and T is a toss where the manner of tossing is heads-inducing. The coin is incapable of landing tails on a trial of that kind.

The specification of a sample space and a kind of trial provides us with a CD -predicate. But a CD -predicate is not yet a chance predicate. Chance predicates are also placeholding predicates of an explicit kind carrying allusion to a background commentary. Part of the commentary is the specification of a CD -predicate. But two distinct chance predicates can presuppose the same sample space. Coin a may have landing heads and landing tails as possible outcomes of a toss and so may coin b . Yet, the chance of heads on a toss of a might be .4 and the chance of heads on a toss of b might be .6. In the next chapter, I address this difference.

† Standard names for elements of $\mathcal{F}(\Omega)$ serve as subscripts for predicates describing outcomes of trials of kind S . $R_{\alpha \cup \beta} = R_\alpha \vee R_\beta$, where $\bar{\alpha}$ indexes $\sim R_\alpha = R_{\bar{\alpha}}$. $\cup_{i=1}^n \alpha_i$ indexes a predicate true of an event if and only if for some α_i , R_{α_i} is true of that event. How sets in $\mathcal{F}(\Omega)$ correlate with the predicates they index is supposed to be determined by the gloss.

Substitution of extensionally equivalent predicates for one another in CD -predicates does not, in general, preserve the extension of the CD predicate. However, if $\alpha = \beta$, then R_α is the same predicate as R_β .

By a *sample space* I mean either the set Ω from which the indices are generated or the set of predicates indexed by the unit sets whose elements are in Ω .

12.1 Chance and Credence

Statistical statements share much in common with disposition statements and statements of ability. To assert that the chance of coin a landing heads on a toss is .5 and the chance of a landing tails on a toss is .5 is to attribute a certain property or condition to coin a just as the attribution to a of the disposition to land heads or tails on a toss or the attribution to a of the ability to land heads on a toss is to predicate a property or condition of a .

Indeed, when the chance predicate is asserted true of a , it is presupposed that a has both the ability to land heads on a toss and the ability to land tails as well as the disposition to land heads or tails and the inability to land both heads and tails.

Furthermore, the chance predicate resembles explicit N -predicates and P -predicates in appearing to have a complex structure. Yet, I have advocated treating N -predicates and P -predicates as though such structure contributes nothing to their semantics. Instead, I propose treating that structure as an abbreviated commentary supplying background information.

Finally, just as knowledge of dispositions and abilities can, under the right conditions, furnish a warrant and backing for appraisals of hypotheses about the outcomes of experiments with respect to serious possibility, knowledge of chances can, again under the right conditions, supply backing for appraisals of hypotheses about the outcomes of experiments with respect to credal probability.

The pseudo structure of chance predicates which quite genuinely supplies background commentary identifies three types of information: (1) a kind of trial S , (2) a sample space Ω indexing possible outcomes of trials of kind S , and (3) a probability measure $F(\beta; \alpha)$ defined for $\beta, \alpha \in \mathcal{F}(\Omega)$ where $\alpha \neq \emptyset$.

Thus, we might write a chance predicate in the form, "is a $C(\Omega/S/F)$."

The background commentary furnished by the specification of a kind of trial and a sample space in the case of a chance predicate corresponds to the commentary afforded by the predicate "is a $CD(\Omega/S)$ " of section 11.10. These elements of the commentary allow the chance predicate to function in appraisals of hypotheses about the outcomes of tests with respect to serious possibility in the manner in which disposition and ability predicates do.

Suppose X knows the following bits of information:

- (i) The chance of coin a landing heads on a toss is .5 and of landing tails is also .5.
- (ii) Coin a is tossed at t .
- (iii) The toss of a at t is also of kind T .

The background commentary informs us that X is committed to the following:

- (iv) Coin a is able to land heads on a toss.

Provided that X 's corpus does not contain information to the effect that a is incapable of landing heads on a toss of kind T and the information about the toss specified in (iii) is all the information about the toss available to X in his corpus, X is obliged to appraise the truth of the hypothesis that the coin lands heads at t as a serious possibility.

Furthermore, X is obliged to rule out the hypothesis that the coin lands both heads and tails as a serious possibility because he knows that a is incapable of landing both heads and tails on a toss.

Finally, X rules out the hypothesis that the coin fails to land either heads or tails because he knows that it is compelled to do one or the other on a toss.

These are the appraisals with respect to serious possibility which X is warranted in making in virtue of his knowledge of chances. They correspond to the appraisals with respect to serious possibility warranted by the associated CD -predicate.

But chance predicates are introduced to do more than that. Two chance predicates may specify the same kind of trial and sample space and yet differ in the chance distribution they

introduce. The significance of this difference is that knowledge of the one chance predicate will warrant different judgments of credal probability for hypotheses about the outcome of experiment than the other.

Thus, knowledge of (i), (ii), (iii), and that the information that the toss is of kind T is stochastically irrelevant (in a sense shortly to be explained) with respect to warrants assigning the hypotheses that the coin a lands heads at t and that the coin a lands tails at t equal Q -values of .5.

This assignment of Q -values is mandated by a restriction on credal states specified in the background commentary for the chance predicate abbreviated by the chance distribution $F(\beta; \alpha)$, which, in our example, states that the chance of heads and the chance of tails on a toss are each .5.

In effect, the specification of a chance distribution imposes a constraint on confirmational commitments by stipulating how knowledge of the specific chance predicate under consideration determines credal states for hypotheses about test behavior.

Strictly speaking, corresponding to each chance predicate, there is a correlated context-independent constraint on confirmational commitments and, in this sense, each chance predicate is associated with its own principle of inductive logic. However, all such principles instantiate the same schema. For this reason, it is convenient to speak of a single principle of inductive logic—the principle of direct inference.

The first task to be faced in elaborating on the concept of chance or objective probability advocated in this book is to offer a formulation of this principle of direct inference.

12.2 Simple Chance Predicates

Suppose a specification of a kind of trial S and a sample space is provided by a chance predicate. In addition, the specification of the chance distribution is given by a function $F(\beta; \alpha)$ such that logic, set theory, mathematics, and standard designators for sets β and α in $\mathcal{F}(\Omega)$ entail that $F(\beta; \alpha) = r$, where r is a standard designator for a real number. In that case, “is a $C(\Omega/S/F)$ ” is a *simple chance predicate* and any singular statement of the form “ a is a $C(\Omega/S/F)$ ” is a *simple statistical hypothesis*.

For example, “has a normal distribution with a mean of 10 and a unit variance” is a simple chance predicate but “has a

normal distribution with mean between -10 and $+10$ ” is not a simple chance predicate. (The latter is customarily called *composite* by statisticians.)

Similarly, “has a chance of heads on a toss equal to the percentage of heroin addicts in New York and a chance of tails equal to the percentage of New York residents who are not addicts” is also composite.

More complicated chance predicates can also be constructed. A measuring instrument may be held to give readings on any object measured normally distributed around the value of the magnitude measured and with unit variance. In effect, the measuring instrument is taken to have a nonstatistical disposition to have chance properties of a certain sort when hooked up to an object to be measured. In this case, the measuring apparatus has the simple chance property of yielding errors in readings on trials on a fixed object being measured which are normally distributed with 0 mean and unit variance.

These and other complications are often present in applications of chance predicates in inquiry. However, the problems of fundamental concern pertaining to direct inference are, in the first instance, questions about how knowledge of simple chance hypotheses justifies or determines credal states.

12.3 Stochastic Irrelevance

Suppose chance setup b is an urn containing 100 coins, half of which have a .4 chance of landing heads on a toss and half of which have a .6 chance of landing heads on a toss. Let a trial of kind S be selecting a coin from the urn blindfolded after the contents of the urn have been thoroughly mixed, so that the method of selection is a “random” one—i.e., one where each item in the urn has an equal chance of selection as any other one—followed by a tossing of the coin selected.

The sample space consists of two possible outcomes—i.e., urn b is capable of responding in two ways to a trial of kind S . The coin selected can land heads or it can land tails.

The chance distribution over this sample space is quite definite. The chance of heads is .5 and the chance of tails is also .5.

Consider now a trial of kind S which is also one where the coin selected is a .4 coin. Call a trial of such a kind a trial of kind $S \& T$.

The sample space is the same as it was for a trial of kind S . So is the chance setup. Both trials are trials on the urn.

But the chance distribution is different for trials of kind S & T . The chance of heads is .4 and the chance of tails is .6.

This is an illustration of *stochastic relevance* of information about a trial on a chance setup. Some trials of kind S on setup b might be of kind T and others not. If the chance distribution over the sample space is different on a trial which is both S and T on setup b from what it is on a trial on b which is S , the information that a trial on b which is of kind S is also of kind T is stochastically relevant to events represented by the sample space Ω for trials of kind S on b . The extra information is *stochastically irrelevant* if it is not stochastically relevant.

Whether such extra information is stochastically relevant or irrelevant depends on the stochastic properties of the coin and its abilities and inabilities. It does not depend on our knowledge.

There is an analogy of limited interest between stochastic relevance and irrelevance and the much less important notion of *modal relevance and irrelevance*. An object b might be able to respond in manner R on a trial of kind S but be quite incapable of responding in manner R on a trial of kind S which is also T (see section 9.8). In that case, the extra information that a trial of kind S on b is also of kind T is modally relevant to the "space" consisting of the "outcomes" of R 'ing and failing to R .

Modal relevance and irrelevance, like stochastic relevance and irrelevance, is a nonepistemological notion. Whether extra information about a kind of trial is stochastically or modally relevant does not depend on X 's knowledge but on the abilities and stochastic properties of the chance setup.

**12.4
Direct
Inference**

Consider some simple chance predicate "is a $C(\Omega/S/F)$." The background supplied by the abbreviated commentary informs us that whatever is a $C(\Omega/S/F)$ is a $CD(\Omega/S)$. What we need now is some account of that part of the background which indicates how knowledge that a chance setup is a $C(\Omega/S/F)$ controls X 's credal state for test behavior.

Consider a corpus K meeting the following conditions:

I. *Knowledge of chance*. K contains the following:

(i) a is a $C(\Omega/S/F)$.

II. *Total Trial Information*. K contains the following information about a trial on chance set up a at t :

(ii) a is subject to a trial of kind S & T at time t .

K contains no additional information concerning the trial of kind S & T at t (except logical consequences).

III. *Stochastic Irrelevance*. K contains the following:

(iii) a is a $C(\Omega/S \& T/F)$.

Let $\alpha, \beta \in \mathcal{F}(\Omega)$ for the sample space Ω associated with the simple chance predicate under consideration. The sentence e_α asserts that the trial of kind S & T occurring at t is followed by an event of kind R_α , and e_β asserts that the trial is followed by an event of kind R_β .

The background commentary supplied by an explicit simple chance predicate stipulates that for any designator a of a chance setup, if the corpus K satisfies conditions I, II, and III, every permissible Q -function in $C(K)$ should satisfy

$$Q(e_\beta; e_\alpha) = F(\beta; \alpha).$$

This stipulation furnishes the application of the principle of direct inference pertinent to the particular simple chance predicate under consideration. Each such application is a principle regulating the use of knowledge of the chance property characterized by the chance predicate involved in determining credal states. For each simple chance predicate, therefore, we have a special principle obligatory on the confirmational commitments of all rational agents who use the predicate. Each such principle is incorrigible in a sense similar to that in which the "reduction sentence" specifying a necessary condition for the application of a disposition or N -predicate is incorrigible.

In section 11.5, it was explained that the background commentary generated by an explicit N -predicate and represented by (5) and (6) of section 11.4 cannot be revised without removing the explicit N -predicate from its status as an explicit predicate of compulsion or disposition.

The same is true, *mutatis mutandis*, of the constraint im-

posed on confirmational commitments associated with simple chance predicates. To be sure, a constraint on a confirmational commitment is not incorrigibly true; for neither Q -functions in credal states nor confirmational commitments have truth values. But such constraints qualify as principles of inductive logic.

On the view I am proposing, there is a separate such principle for each simple chance predicate. However, we can consider a schema to characterize the entire family of such constraints. It is such a schema which shall be referred to as the *principle of direct inference*. Strictly speaking, this principle is not itself a principle of inductive logic; rather, it characterizes a family of such principles. But it will do little harm, at least in most contexts, to speak as though it itself is a single principle.

12.5 The Semantics of Chance

According to my proposal, a simple chance predicate is to be regarded as primitive for the purpose of supplying necessary and sufficient satisfaction conditions with relatively incorrigible status. To accord chance predicates such "theoretical" status may not, perhaps, appear objectionable—except for one thing. It is sometimes held that an adequate "interpretation" of chance or of objective or statistical probability should supply a model of the formal calculus of probability. In failing to do this, my proposal violates an important condition of adequacy.

I reject the condition of adequacy; but in order to meet the challenge of the objection, let me repeat once more what I am disclaiming when I refuse to supply a semantics for simple chance predicates.

Consider first the case of N -predicates of the type "is an $N(R/S)$." Syntactic rules may be formed allowing for the formation of such a predicate from well-formed predicates or open sentences R and S . When a semantics is sought for such predicates, the usual expectation is that satisfaction conditions for "is an $N(R/S)$ " will be controlled by the semantical interpretation of R and S .

Most authors would agree that expressions such as R and S occur nonextensionally in predicates like "is disposed to R when S 'd" and, hence, would not undertake to determine the extensions of N -predicates as functions of the extensions of

their constituents. Students of the voluminous literature initiated by G. Frege and B. Russell will no doubt entertain various moves for spelling out satisfaction conditions for N -predicates in terms of the meanings (somehow specified) of their parts.

Such efforts serve no useful purpose, as far as I can see, in assisting us to come to grips with the role of knowledge of dispositions, incapacities, and capacities in inquiry and deliberation. I have suggested, therefore, construing the "constituents" as providing an abbreviated side commentary specifying the truth and lawlikeness of appropriate reduction sentences. In the case of N -predicates, the background commentary supplies necessary satisfaction conditions in terms of satisfaction conditions for the constituent predicates R and S ; but that is all.

Chance predicates offer a more elaborate set of constituents. There are descriptions of a kind of trial and of kinds of outcomes indexed for formal convenience by sets in a field generated by a sample space Ω . The predicates characterizing these kinds of events occur in nonextensional contexts in the chance predicates; and, for substantially the same reasons as in the case of N -predicates, are to be treated in the same way. The marginal commentaries for simple chance predicates supply necessary satisfaction conditions for these predicates in terms of descriptions of test behavior.

The remaining "constituent" of the chance predicate is the probability function $F(\beta; \alpha)$ defined over sets belonging to $\mathcal{F}(\Omega)$. The background commentary, in this case, does not supply any satisfaction condition at all for the chance predicate. Instead, an epistemological constraint is imposed in the guise of a condition on confirmational commitments.

Thus, my refusal to supply a semantics for simple chance predicates means at least this much: necessary and sufficient satisfaction conditions for chance predicates are not determined as functions of the semantics of their constituents.

But my indifference to semantics goes still further.

Suppose that " a is a $CD(\Omega/S)$ " is true.

Someone might seek necessary and sufficient truth conditions for " a is a $C(\Omega/S/F)$ " conditional on the truth of the supposition. If this could be done, one might claim to have offered some sort of interpretation of the probability function

and, in this sense, an interpretation of the calculus of probabilities.

I do not know how to provide such an interpretation, and think efforts to do so are both gratuitous and diversionary as far as the main problems about chance are concerned—which problems are to furnish an account of how knowledge of chance controls credal judgments (and other judgments) about test behavior.

The project of interpretation might be weakened yet further by seeking some necessary truth condition for the simple chance statement on the supposition that the *CD*-statement is true.

This latter project is, I suspect, the one which J. Venn, R. Von Mises, and C. S. Peirce undertook when they sought “interpretations” of the calculus of probabilities by reference to limits of relative frequencies in sequences of trials of the sort Von Mises called “collectives.” On this view, “*a* is a $C(\Omega/S/F)$ ” is true only if *a* is disposed to yield results of kind R_α with a limiting relative frequency $F(\alpha)$ (where $F(\alpha) = F(\alpha; \Omega)$) in an infinite sequence of repetitions of trials of kind *S* satisfying the conditions for a collective.†

I find it difficult to take the infinitely long run as much more than a bad joke. If coin *a* has a .5 chance of landing heads and a .5 chance of landing tails on a toss, it is able to land heads on a toss and is also able to land tails. Under the circumstances, it seems utterly implausible to claim that coin

† The truth is that neither Venn, Peirce, nor Von Mises are very clear. The dominant theme in Von Mises is that the formal calculus of probability needs an interpretation which he provides by introducing the notion of a collective and limits of relative frequency in collectives. Von Mises claims that the probability calculus so interpreted is a scientific theory.

Prima facie Von Mises is committed to the reality of such collectives. However, he denies this. He claims collectives are idealizations as are, on his view, all theoretical concepts. (See Von Mises, *Probability, Statistics, and Truth*, London: Allen and Unwin, 1957, pp. 6–8.)

The best I can do in order to make sense of these mysteries is to appeal to an analogy with measuring length. If we claim that a rod is exactly 1.89 meters long, we do not claim that measurement will yield this value as an exact reading. Setting aside random error, measuring rods yield readings with some round-off error so that we can obtain values only to the nearest value at some decimal place.

Nonetheless, we might attribute a length of 1.89 meters to the rod intending to claim that the rod is disposed to furnish interval-valued readings containing the value 1.89 (setting aside random error) no matter how narrow the interval

a is incapable of landing heads every time in a million tosses or even an infinite sequence of tosses. Von Mises agrees about the million tosses but it is crucial to his interpretation of chance that coin *a* be incapable of landing heads every time in an infinite number of tosses.

To be sure, on a trial of coin *a* consisting of an infinite sequence of tosses where the chance of *r* heads in *n* tosses for an initial segment of *n* tosses of the infinite sequence is $\binom{n}{r}(.5)^n$, the chance of all heads in the infinite sequence is 0. Given that coin *a* has the property of yielding heads on a single toss and tails on a single toss with an equal chance of .5 if and only if it has the chance property just described, it might be thought that coin *a* is incapable of landing heads every time in an infinite sequence of tosses because the chance is 0.

Observe, however, that the chance of heads and tails in any specific order leading to a limit of relative frequency equal to .5 on the infinite sequence of tosses is also equal to 0. If the coin is incapable of landing heads every time in an infinite sequence of tosses because the chance of doing so is 0, the chance of landing heads and tails in any particular sequence should also be impossible for the coin to achieve on an infinite sequence of tosses for the same reason even if the limit of relative frequency is .5.

Intuitions about the infinitely long run, however, are notoriously shaky. Neither the judgments of limit of relative-frequency theorists nor those who share my attitude will settle the dispute.

is—i.e., how precise the measuring instrument is. (I ignore complications due to quantum mechanics.)

Von Mises seems to think that when a chance of .5 of landing heads on a toss is attributed to coin *a*, this is an idealization in the sense that “measurement” of this chance property by means of relative frequency of heads in a finite albeit large number of tosses will not invariably yield a relative frequency equal to .5. However, increasing the number of tosses (the alleged analogue to making more precise measurements) will in the long run lead to relative frequencies converging on .5.

Understood in this manner, Von Mises is not implying that chance statements presuppose that measurements are ever made. Not even an initial segment of a collective need occur. In effect, the chance setup is alleged to have a disposition to a limit of relative frequency of .5 in a sequence of tosses forming a collective analogous to the disposition attributed to the rod to respond in certain ways to measurement. This is the construal of Von Mises’ view I am now using. I am not sure it is accurate to his intention; but it is the most charitable reading I can think of.

If the proposed limit of relative-frequency interpretation of chance could shed light on connections between chances and outcomes of finite sequences of trials, including very long sequences, the objection just raised would scarcely be decisive.

Von Mises' interpretation of the calculus of probability fails to supply an understanding of these connections. To supply an understanding of these connections requires appeal to some epistemological constraint. I contend such constraints come, in the first instance, in the form of principles of direct inference. Once we have such principles, we do not need Von Mises' interpretation of the calculus of probability to obtain an adequate understanding of the concept of chance.

My chief objection to Von Mises' view, therefore, is that by itself it contributes nothing to the understanding of chance and, what is worse, generates a diversion from the central problems by directing attention to the irrelevant question of how to define a collective.

Von Mises proclaimed that "the theory of probability is a science similar to others."¹ His idea seems to have been that the "theory of probability" is the product of furnishing the formal calculus of probability with a model or interpretation which could be used in explaining and predicting "mass phenomena" in a wide variety of domains.

The idea is obviously untenable. Chance predicates occur in the formulation of laws and theories covering diverse subject matters, both minute and large. To suppose that there is a single science of chance covering quantum mechanics, genetics, statistical mechanics, and various branches of sociology and economics is reminiscent of the grandiose and Pickwickian claims sometimes made on behalf of general systems theory and general semantics.

One of the advantages of taking chance predicates to be primitive placeholding predicates is that this view emphasizes the absurdity of a science of probability. The idea is as bizarre as the idea of developing a science of abilities, incapacities, dispositions, and possibilities. (Modal realists have not been as frank about their ambitions as Von Mises was about his. But it may be worthwhile registering an anticipatory protest just in case someone harbors in his breast the ambition of

devising a science of possible worlds that would be as acceptable as any other science.)

12.6 Chance and Coherence

There is one demand made by those who seek a semantics for chance which has thus far not been considered.

The function $F(\beta; \alpha)$ has been required to be a probability measure over the domain $\mathcal{F}(\Omega)$. No rationale for this requirement has been offered within the framework of my account of chance. Why should the function F obey the requirements of the calculus of probabilities?

Recall that the point of introducing simple chance predicates into the language L is to have a means of describing conditions knowledge of the presence of which warrants (everything else being equal) credal states for hypotheses about outcomes of trials known to have occurred. According to the principle of credal coherence, the permissible Q -functions representing such credal states should obey the calculus of probabilities. The principle of direct inference associated with a simple chance predicate stipulates that the permissible Q -function for hypotheses about the outcome of a given trial should be identical with the function $F(\beta; \alpha)$ given the appropriate correlation of "events" β and α with hypotheses about the outcome of the trial.

The rationale for credal coherence is found in the account of how permissible Q -functions function in deliberation and inquiry in the evaluation of feasible options with respect to expected utility. The principle of direct inference is the keystone to the account being proposed for characterizing the way knowledge of chances determines credence. To guarantee the consistency of these two principles of inductive logic, the chance distribution for any simple chance predicate should obey the requirements of the calculus of probabilities.

I see no reason for searching for some rationale additional to this.

12.7 Chance and Credal Consistency

Consider the following pair of sentences:

- (1) The chance of coin a landing heads on a toss is .5 and the chance of tails is also .5.
- (2) The chance of coin a landing heads on a toss is .1 and of landing tails is .9.

According to my proposals, (1) and (2) predicate distinct predicates true of the chance setup a . There is no logical inconsistency in doing so. More to the point, nothing has been assumed which precludes the joint truth of the two assumptions.

We can easily introduce as an additional gloss in the background commentary of every simple chance predicate a proviso that no simple chance predicate specifying the same kind of trial and sample space but a different chance distribution is true of any chance setup of which the chance predicate in question is true. I favor a slightly weaker stipulation precluding the joint truth of such predicates only when the setup is capable of being subjected to the kind of trial under consideration.

We still lack, however, a rationale for introducing such a proviso.

Suppose X 's corpus K contained (1) and (2) and, in addition, information to the effect that coin a is tossed at time t . Suppose X knows that all additional information about the toss is stochastically irrelevant concerning the particular toss at t .

The principle of direct inference mandates assigning the hypotheses that the coin lands heads up and that it lands tails up equal Q -values of .5. The principle of direct inference also mandates assigning the two hypotheses the Q -values of .1 and .9 respectively.

That is to say, all permissible Q -functions are supposed to make both sorts of assignments. No permissible Q -function can do so. Hence, the set of Q -functions will be empty—in violation of credal consistency.

The only way to escape the violation of credal consistency and preserve the principle of direct inference is to guarantee that the corpus K is inconsistent. It will be so if the background information is guaranteed to contain the provisos preventing the joint truth of (1) and (2) mentioned above.

Consider the following pair of statements:

- (3) The chance of die b landing with i spots up (for i from 1 to 6) is $1/6$.
- (4) The chance of die b landing with one spot up is $1/2$ and with some other number of spots up is $1/2$.

In the case of this pair of statements, the sample space indexing possible outcomes is different; hence, the example is not the same as the one where (1) and (2) are considered. However, some possible outcomes generated by one of the sample spaces coincide with those generated by the other and, for those outcomes, the chances differ.

Considerations similar to those adduced previously argue for inclusion in the incorrigible background assumptions of stipulations to the effect that (3) and (4) are not both true.

Refer back to the urn b of section 12.3. Compare the following statements:

- (5) The chance of heads up on a trial of kind S & T is .4.
- (6) The chance of heads up conditional on obtaining a .4 coin on a trial of kind S is .4.

Both statements predicate a stochastic property of urn b . The predicates are distinct and are primitive. There is no ground in what has been explicitly stated thus far for supposing that the first predicate should be true of exactly the same setups as the second. Yet, it should be apparent that they should.

The reason is this. If X knows that (6) is true and that a trial of kind S has occurred and, moreover, lacks additional stochastically relevant information about the trial, all Q -functions in his credal state should assign a Q -value of .4 to the hypothesis that the coin selected lands heads up conditional on the coin selected being a .4 coin.

If, at that stage, X should add to his corpus the information that the coin selected is a .4 coin, confirmational conditionalization mandates that X assign an unconditional Q -value of .4 to the hypothesis that the coin selected lands heads up.

But this new corpus is one where X knows that the trial is of kind S & T . If the chance of heads up on a trial of kind S & T were assumed to be different than .4, direct inference would mandate a Q -value different from .4 to heads up. Hence, relative to the same corpus, X would be obliged to both assign and refuse to assign a unique Q -value of .4 to the hypothesis that the coin lands heads up. To avoid this, the chance specified in (5) should be the same as that specified in (6).

I shall not undertake the effort of specifying in a systematic

and exhaustive manner the background postulates for chance predicates which could be introduced into the urcorpus in the manner indicated. The principle operative here for justifying such postulates should, however, be fairly clear. The principle of direct inference should apply without conflicting with principles of credal coherence and consistency.

12.8
The "Existence" of
Chance

All setups which are $C(\Omega/S/F)$'s are $CD(\Omega/S)$'s. In general, the converse does not hold. The question arises, however: Is every $CD(\Omega/S)$ a $C(\Omega/S/F)$ for some function $F(\beta; \alpha)$ or other?

Suppose box a has two compartments. The left compartment contains 40 black balls and 60 white balls while the right compartment contains 40 red balls and 60 blue balls. A trial of kind S is selecting a ball at random from the left compartment and a trial of kind S' is selecting a ball at random from the right compartment. The sample space for trials of kind S consists of drawing a black ball and drawing a white ball. The sample space for trials of kind S' consists of drawing a red ball and drawing a blue ball. Chances are defined for both kinds of trials over their respective sample spaces.

Consider trials of kind $S \vee S'$. There is indeed a sample space consisting of drawing a red ball, a blue ball, a black ball, and a white ball. However, there is no chance distribution over the sample space.

To see why no chance distribution is defined, consider that the sample space for trials of kind $S \vee S'$ is such that a result consisting of obtaining a red or a blue ball is equivalent to obtaining a result of conducting a trial of kind S . Similarly obtaining a black or white ball is equivalent to obtaining a result of conducting a trial of kind S' . If there were a chance distribution over the sample space for trials of kind $S \vee S'$, that distribution would assign a chance to obtaining a result of a trial of kind S on a trial of kind $S \vee S'$. (Similarly for results of a trial of kind S' .) Thus, conducting a trial of kind $S \vee S'$ would be conducting a trial of kind S with some definite chance or statistical probability.

There is no a priori consideration precluding such chances; but there is no guarantee that such chances are defined either. In the example under consideration, we would normally deny that they are.

To be sure, if we were to conduct a trial consisting of tossing

a coin with a known chance of heads and of tails and then conduct a trial of coin S if the coin lands heads and a trial of kind S' if the coin lands tails, there would be a definite chance of conducting a trial of kind S (and of kind S') on such a trial. But the trial is not merely of kind $S \vee S'$. It has an additional feature—namely that whether a trial of kind S or of kind S' is conducted depends on the outcome of a stochastic process. The trial is of kind T where all trials of kind T are trials of kind $S \vee S'$, but not conversely.

Suppose that chance setup a is a $C(\Omega/S/F)$ for some chance distribution F . Consider trials of kind S & T . a will be a $CD(\Omega'/S \& T)$ for some $\Omega' \subseteq \Omega$. I assume that for some F' defined over $\mathcal{F}(\Omega')$, a is a $C(\Omega'/S \& T/F')$. On the other hand, as the previous example illustrates, if chances are defined on a trial of kind S & T for setup a , they need not be defined for the same setup relative to trials of kind S .

Thus, given that the chance of coin a landing heads on a toss is p and of landing tails is $1 - p$, there is some value $0 \leq p^* \leq 1$ such that the chance of coin a landing heads on a toss by Isaac Levi on June 30, 1978, at 12:00 noon in 715 Philosophy Hall is equal to p^* .

Suppose X knows that the extra information about the toss on June 30, 1978, at 12:00 noon in 715 Philosophy Hall additional to its being a toss of coin a is

- (i) that it took place on June 30, 1978;
- (ii) that it took place in 715 Philosophy Hall;
- (iii) that the toss was conducted by Isaac Levi.

If, in addition to this, X knows that this information is stochastically irrelevant concerning outcomes of tossing coin a —i.e., that $p^* = p$, then X 's credal state for hypotheses concerning the outcome of the toss will depend on his knowledge of the value of p and, if he is in suspense concerning alternative hypotheses concerning that value, on his credal state over these alternatives.

In particular, if e_H is the hypothesis that the coin lands heads and h_p specifies that the value of the chance of heads on a toss is p , X 's credal state must be such that $Q(e_H; h_p) = p$.

This much is mandated by direct inference.

Thus, X has the right to ignore the information contained in (i), (ii), and (iii) provided he knows that this information is stochastically irrelevant. If X lacks such knowledge, he must base his direct inference on whatever knowledge he does have concerning the chances of heads on a trial having the features specified by (i), (ii), and (iii).

Of course, X can and, indeed, must ignore information he does not have in his corpus. For example, X may know that tosses of a certain kind are invariably heads-inducing and that tosses made under other conditions are invariably tails-inducing so that information that a toss is heads-inducing is stochastically relevant. But X does not know of the trial under consideration whether it is heads-inducing or tails-inducing and, hence, can ignore these prospects.

Thus, the credal states determined via direct inference depend critically on X 's knowledge of stochastic irrelevance. The point is not entirely without precedent. Whether one can infer from knowledge that a is able to R on a trial of kind S and that a is subjected to a trial of kind S an appraisal of the hypothesis that responds in manner R with respect to serious possibility depends on the extra information one has concerning the trial in question.

One good reason for allowing chances to be defined over sample spaces relative to trials of a highly specific kind now emerges.

Suppose X knows, as before, that coin a is to be tossed in 715 Philosophy Hall at the time mentioned and by Isaac Levi. Suppose X also knows that the chance of coin a landing heads on a toss is p .

Unlike the previous case, however, we do not allow chances to be defined relative to kinds of trials specified to be tosses meeting conditions (i), (ii), and (iii). I suspect that many authors would be inclined to take such a position—in opposition to my view that if chances are defined relative to trials of kind S they are defined relative to trials of kind $S \& T$.

One implication of such a view is that information furnished by (i), (ii), and (iii) no longer qualifies as stochastically relevant or irrelevant. It is stochastically nonrelevant.

But saying this does not decide whether the extra information precludes or does not preclude X 's adopting the credal state where all Q -functions are such that $Q(e_H; h_p) = p$. In

other words, nothing is decided as to whether X may or may not ignore the extra information about the toss in making direct inferences.

If one says that when chances are undefined for trials described in a manner which includes the extra information that the extra information may be ignored, one is treating the extra information as if it were stochastically irrelevant. However, by refusing to allow assumptions to that effect into the corpus, one prevents recognition of the feasibility, on some occasions, of revising such judgments of stochastic irrelevance.

If, on the other hand, a blanket license to ignore stochastically nonrelevant information is not given in direct inference, our principle of direct inference is utterly deprived of serious applicability; for few realistic situations would arise where there would be no such stochastically nonrelevant information.

By insisting that chances are defined relative to trials of kind $S \& T$ if they are defined relative to trials of kind S , the conditions under which extra information about trials of kind S may be ignored can be formulated in a straightforward manner.

12.9 More on the Existence of Chance

The various postulates stipulated to belong "in the background" for chance predicates, the specifications of the scope of definability of chance, and the characterization of principles of direct inference have all been motivated by a view of the role of chance predicates as placeholder characterizations of properties of objects knowledge of which can be used to evaluate hypotheses about test behavior with respect to credal probability.

Thus, not only are chance predicates placeholders (as are disposition and ability predicates) but their placeholder function can be explicated only by appealing to restrictions on confirmational commitments. In a sense, the characterization of chance proposed here is more intimately bound up with epistemological considerations than is the account of disposition predicates.

Nothing in the view proposed, however, implies that chance predicates are interpreted epistemologically. Chance statements do not describe cognitive states of agents. They bear truth values and attribute properties or conditions to objects

which they have or lack independent of the subjective states of knowing subjects. Agents can accept or remove chance statements from their corpora and may suspend judgment as to their truth and they may adopt credal states regarding chance hypotheses.

The epistemological component of the analysis proposed here is designed to identify an important connection between knowledge of chance and credal judgments concerning test behavior. In lieu of reduction sentences which establish such connections for N -predicates, principles of direct inference functioning as principles of inductive logic are employed.

But if chance is parasitic on credal probability in this way, why should chance predicates be introduced into the language altogether? Could we not dispense with chance statements and rest content with credal judgments about test behavior?

This sort of conclusion has been advanced by B. De Finetti, L. J. Savage, and others who have undertaken to show that hypotheses about objective chance are expendable in scientific inquiry and in practical deliberation.²

But chance predicates are no more expendable in inquiry and deliberation than disposition predicates are.

Some iron bars attract iron filings placed near them and others do not. As a first step toward understanding the differences between the two sorts of iron bars, X may say that one sort of bar has a disposition to attract iron filings and the other does not. Of course, this description of the difference is but a first step. That is why explicit disposition predicates are placeholders for more adequate characterizations of the relevant differences. Nonetheless, they have an important function, and in many instances an indispensable one, in inquiry and deliberation.

Suppose X 's credal state for the hypothesis that coin a will land heads r times in n tosses conditional on its being tossed n times at time t is $\binom{n}{r}(.5)^n$. His credal state for the hypothesis that coin b will land heads r times in n tosses conditional on its being tossed n times at t is $\binom{n}{r}(.9)^r(.1)^{n-r}$.

As far as the requirements of credal coherence are concerned, X may endorse a state embodying both credal commitments even if he assumed that a and b are similar in all respects except that a came off the mint just before b did.

Yet, not even De Finetti would say under these circum-

stances that such a credal state makes sense. There is nothing in De Finetti's view of credal rationality which prevents such a credal state from being adopted; but there would be wide agreement that such a credal state makes no sense unless there is some significant difference in the characteristics of coin a and coin b .

That is not to say that X should be in a position to offer an explanatorily adequate characterization of the difference between the coins; but he should be committed to the view that there is a difference in traits. The coin a has some property C such that given knowledge that an object has C , *ceteris paribus*, X 's credal state for hypotheses specifying relative frequencies of heads on n tosses should be as specified above. Similarly, b has some property C' knowledge of the presence of which licenses a credal state of the sort attributed to hypotheses about b 's behavior.

One way of putting it is to say that coin a is unbiased whereas coin b is heavily biased in favor of heads. Another way to put it is to specify the explicit simple chance predicates which are true of a and of b concerning outcomes of n tosses.

De Finetti is quite right when he complains of the deficiencies of descriptions of differences between coin a and coin b in terms of differences in chances or objective probabilities. He is wrong, however, in supposing that the deficiencies are to be removed by restricting credal judgments to hypotheses about test behavior and forbidding the acceptance of chance hypotheses into evidence or the assignment of credal probabilities to them.

The defects in chance predicates are to be removed not by eliminating chance statements from the language because chance is alleged to be metaphysical moonshine, but by integrating chance predicates into theories through inquiry as is attempted in genetics, statistical mechanics, and quantum mechanics in differing ways.

Those who are unhappy about treating chance predicates as primitives for the purposes of semantics may also be granted their sop to Cerberus. *Qua* placeholders chance predicates have deficiencies for purposes of explanation and prediction and these deficiencies may be removed only by integrating the chance predicates into explanatorily satisfactory theories. One way (but not the only one) in which this might be done is

through finding theoretically adequate predicates extensionally equivalent to the chance predicates in question.

This point leads to another partial concession to De Finetti. Suppose that the chance predicate "has a .5 chance of landing heads on a toss and a .5 chance of landing tails on a toss" is discovered to be extensionally equivalent to a description of the physical characteristics of the coin—say in terms of its shape and the distribution of mass over the coin; and this equivalence is taken to hold as a matter of law.

If this physical description is abbreviated by using the predicate "is an M ," we are entitled to make direct inferences from knowledge that coin a is an M and that it has been tossed to the judgment that the hypothesis "coin a lands heads on that toss" has a degree of credence equal to .5.

This direct inference, however, is not licensed by principles of inductive logic alone. It is only because the corpus contains the information that "is an M " is true of all and only those objects having an equal chance of landing heads on a toss and of landing tails on a toss that we are entitled to make credal evaluations via direct inference in this fashion.

Thus, in a certain sense, the placeholder chance predicate is expendable for purposes of direct inference; but only in a sense. The replacement of the placeholder simple chance predicate does not enable us to dispense with reasoning from knowledge that the setup has the physical property M to credal judgments about test behavior. And our warrant for such direct inference depends on knowledge of the equivalence in extension of the physical predicate and the chance predicate it replaces.

Thus, the progress which has been made in finding a physicochemical basis for genetic processes has not allowed us to dispense with direct inference in making predictions about the traits of offspring of parents of known genotype.

Some care should be taken to avoid confusing this last remark with a commitment to some strong form of determinism. Replacing the chance predicate for the coin by the extensionally equivalent "is an M " will not and cannot yield a deterministic law to the effect that all M 's land heads when tossed. On the other hand, it may turn out that some tosses meet additional conditions such that all M 's land heads when tossed in a manner meeting these additional conditions. Fur-

nishing an extensionally equivalent "basis" for a chance predicate in explanatorily acceptable terms does not suffice for the purpose of converting statistical explanations into nonstatistical ones. But it does not preclude finding nonstatistical explanations either.

12.10 Chance and the Long Run

Even those who might grant that chance predicates are placeholders useful in predicting and explaining test behavior in a stopgap fashion may object that the test behavior of interest concerns the relative frequencies with which results of various kinds occur in long sequences. Knowledge of the chances of heads and of tails on a toss of coin a is of interest for the purpose of explaining or predicting the relative frequency with which a lands heads on long sequences of tosses.

Nothing said thus far provides a link between knowledge of chances and test behavior on repeated trials. The principles of direct inference which have been proposed specify, in the case of the coin, how X should evaluate credal probabilities for hypotheses about the outcome of a single toss given knowledge of chances of heads and tails on a single toss. What are the implications of these principles concerning how knowledge of chances of heads and of tails on a single toss control evaluations of hypotheses concerning relative frequencies of heads and of tails in the long run?

Two points may be made by way of a preliminary response.

(1) Even though it is true that knowledge of chances is used to explain and predict relative frequency in long runs, it is also used to evaluate hypotheses about outcomes of single trials.

(2) The account of chance proposed here provides for appraisal of hypotheses concerning relative frequencies of outcomes on sequences of trials of some kind. Suppose that X knows that the chance of coin a landing heads r times and tails $n - r$ times in a definite order on a trial consisting of sequence of n tosses is $(.5)^n$. He then knows that the chance of r heads in n tosses regardless of order is $\binom{n}{r}(.5)^n$. The kind of trial is tossing coin a n times in a row. The sample space consists of the 2^n distinct sequences of heads and tails in n tosses. Given the additional knowledge that on some specific occasion coin a has been tossed n times and that all further information about the n tosses is stochastically irrelevant, X is obligated

by direct inference to assign to each hypothesis concerning the relative frequency of heads a Q -value equal to $\binom{n}{r}(.5)^n$.

This answer is not fully responsive to the demands of the objection. Contrast the two simple statistical predicates:

- (i) "has a .5 chance of landing heads on a toss and a .5 chance of landing tails on a toss."
- (ii) "has a chance of $\binom{n}{r}(.5)^n$ of yielding r heads on n tosses."

Nothing in the account proposed guarantees that the first predicate is true of the same objects as the second and only these. Direct inference establishes a link between knowledge that the second predicate is true of chance setup a and its behavior on n tosses. Knowledge that the first predicate is true of a establishes a basis for judging outcomes of a single toss, but not for $n > 1$.

We might propose introducing another postulate in the ur-corpus specifying that the first chance predicate is true of a coin if and only if the second is (and that this is so for arbitrary n).

Sometimes X does and should accept such an assumption in his corpus. When he does he is assuming that sequences of n tosses are sequences of stochastically independent tosses where the chance of heads remains the same on each toss. (To say that the tosses are stochastically independent is to say that the chance of heads on the i th toss is the same regardless of the outcome of the i' th toss.)

Such assumptions are eminently open to revision and should not be embedded in UK .

Thus, X might believe that coin a is not very durable so that each toss alters the chance of heads on the next toss and that how it alters the chance is a function of the result of the preceding tosses. X might believe that coin a , which has never been tossed, has a .5 chance of landing heads on a toss as long as it remains untossed. Yet, he might not believe that the chance of r heads on n tosses is $\binom{n}{r}(.5)^n$.†

† I have been too glib in speaking of tosses rather than tosses under such and such conditions. X may not believe that the chance of heads on a toss of coin a in a sequence of n tosses is the same for each toss in the sequence or that the tosses are stochastically independent. Yet, he may believe this for tosses under conditions C .

To be sure, if X believes that the chance of the untossed coin's landing heads on a toss is .5 as long as it is untossed, he should also believe that the chance of heads on the first of a sequence of n tosses is .5. This assumption relies on the conviction that the behavior of the coin on tosses after the first in an n -fold sequence does not influence its behavior on the first toss. In principle, the assumption is open to revision; but, given the current state of knowledge, it will, I trust, be widely accepted.

The crucial point is the absence of any conceptual, incorrigible, or a priori links between chances on single trials and chances on sequences of trials of the same kind. The quest for such links is misguided and ought to be abandoned.

Of course, processes involving sequences of stochastically independent trials where chances are constant for each trial in the sequence constitute an important chapter in discussions of chance and statistical inference. But they constitute but one chapter. There are others. Chances on single trials are linked with chances on sequences of trials in other ways as well. The situation for each kind of chance setup and trial has to be investigated separately.

It is also high time to question the old dogma that kinds of trials for which chances are definable should be repeatable. Suppose bottle a is disposed to break into pieces when dropped. The kind of trial (dropping the bottle) is not repeatable on the setup a . Yet the bottle may be said to have the disposition. We can also say that the bottle a has a chance p or breaking into 10 pieces when dropped even though the trial is not repeatable. Illustrations of the nonrepeatability of kinds of trials on the same setup abound in quantum mechanics. The phenomenon, however, is not restricted to that domain.

To be sure, bottles similarly "prepared" and, hence, having the same dispositions or "propensities" to break can be subjected to the test. In that sense, the kind of trial is repeatable.

Observe, however, that the assumption that all the similarly prepared bottles have an equal chance of breaking into 10 pieces when dropped presupposes the intelligibility of talking about the chance of any specific one of the bottles breaking into 10 pieces when dropped.

Moreover, there is no principle which guarantees the repeatability of trials of some kind on similarly prepared setups.

Finally, chances are definable relative to trials which can be conducted on at most one chance setup and at most once—e.g., trials which are tossings of coin *a* by Isaac Levi at 12 noon in Philosophy Hall on June 30, 1978.

In rejecting repeatability as a condition on kinds of trials, I do not mean to deny that the primary importance of chance concepts in inquiry and deliberation relates to relative frequencies on repeated trials. This important and true observation does not warrant the conclusion that there is some incorrigible connection between chance and relative frequency. Chance is sometimes related to relative frequency and sometimes not; and when it is so related, the connections can be of different sorts. Which links obtain depend on the chance setup, its properties, and the kinds of trials being contemplated. Only inquiry can find out what these connections are. No amount of armchair reflection will determine the outcome.

**12.11
Statistical
Prediction**

Reduction sentences establish the only incorrigible links between dispositionality and test behavior. Principles of direct inference determine the only fixed links between knowledge of chances and judgments about test behavior.

However, knowledge of chances is often used to make predictions about relative frequencies of outcomes in sequences of trials. Thus, if coin *a* which is known to have a chance *p* of landing heads on a single toss is tossed *n* times where the tosses are known to be stochastically independent and with constant chance *p* of heads on each toss, it is to be expected that coin *a* will land heads approximately 100*p*% of the time.

What does “it is to be expected” mean here? I suggest that *X* is warranted in expanding his corpus inferentially to accept as strongest the conclusion that the relative frequency falls within some small interval of values around *p*.

Of course, the legitimacy of an inductive expansion depends on the set of potential answers as determined by an ultimate partition, the demands for information as represented by an *M*-function, and the value of an index *q*. But in contexts of the sort under consideration, it is not too difficult to suppose that demands will normally be of a sort where *X* is interested in predicting relative frequencies.

This means that in a binomial process of the sort being envisaged, on *n* trials, there are *n* + 1 hypotheses specifying

distinct relative frequencies of positive outcomes (heads, say). Accepting any one of these is a potential answer and each of these answers should bear equal *M*-value of $1/(n + 1)$.

The ultimate partition can be taken to consist of these *n* + 1 alternatives or, if one likes, of the 2^n hypotheses specifying distinct sequences of heads and tails in the *n* tosses. All sequences yielding *r* heads on *n* tosses—there will be $\binom{n}{r}$ of these—should bear equal *M*-value and, hence, *M*-value equal to $r!(n - r)!/(n + 1)!$. In any case, the hypothesis that *r* out of *n* tosses land heads will be rejected if and only if

$$p^r(1 - p)^{n-r} < \frac{qr!(n - r)!}{(n + 1)!}$$

This holds if and only if

$$\binom{n}{r}p^r(1 - p)^{n-r} < \frac{q}{(n + 1)}$$

It can be shown³ that as *n* increases the warranted conclusion is that the value of *r/n* will be very close to *p*, and that the interval around *p* becomes smaller and smaller as *n* increases provided that *q* is held constant.

If reiteration of the rule is allowed, the stable conclusion (see section 2.8) will be even stronger.

Similar results can be obtained for repeated trials where there are *k* possible outcomes and *n* stochastically independent repetitions with chances the same on every trial. Well-behaved results can also be obtained for other distributions as well (such as repetitions of trials where outcomes are represented by normally distributed random variables).

Of course, these results, which plainly conform to presystematic judgment, depend not only on *X*'s knowledge of chances, but also on his demands for information. Critics can object that bizarre results can also be obtained using criteria for inferential expansion of the sort I have advocated by altering the ultimate partition or the *M*-function.

Remember, however, that the legitimacy of an inferential expansion depends on the demands for information which the agent is committed to realizing. If the agent is seriously concerned with predicting relative frequencies of outcomes on *n* trials, he will be committed to ultimate partitions and *M*-functions yielding results such as those I have just specified.

12.12
Subjunctive
Conditionals

One cannot have it both ways. One cannot complain that my view of inferential expansion cannot account for how knowledge of chances warrants predictions about relative frequencies by citing cases where agents are not interested in such prediction but have other demands for information.

Statements of chance, like disposition statements, allegedly support subjunctive conditionals.

Consider a fair coin a . Asserting that it is fair supports the counterfactual judgment that if a were tossed 1,000 times, it would land heads approximately 500 times. It also supports the counterfactual judgment that were a tossed 1,000 times, it could land heads every time.

Students of possible-world semantics will no doubt worry about inconsistency. They will seek to avoid trouble by forbidding the first appraisal. They will suggest instead that the chance statement supports the judgment that if a were tossed 1,000 times, in all probability it would land heads approximately 500 times. Such a move is high-handed adhocery designed solely to avoid trouble in an outlook built on fantasy in the first place. There is another way out that is no better. It may be suggested that the measure of similarity between this world and other possible worlds used in the first appraisal and the second is different. Clearly this suggestion is as ad hoc as the first one.

The approach to counterfactuals I favor can do much better.

Let the initial corpus K contain the following:

- (i) Coin a has a .5 chance of heads on a toss and a .5 chance of tails; and repeated tosses are stochastically independent and do not alter the chances.
- (ii) Coin a is not tossed during time t .
- (iii) Coin a does not land heads or tails during time t .

Let K_1 be obtained from K by removing (ii) with minimum loss of information. Assuming that (i) is not to be removed and that the background information specifying that a is compelled to land heads or tails on a toss remains intact, removing (ii) compels removal of (iii).

Let K_2 be obtained by adding the hypothesis that coin a is tossed 1,000 times during t to K_1 . Relative to K_2 , the hypoth-

esis that coin a lands heads every time in the 1,000 tosses is a serious possibility. Thus it is that (i) supports the counterfactual "if a were tossed 1,000 times, it could land heads every time."

Consider now an inferential expansion relative to K_2 , seeking to predict the relative frequency of heads on the 1,000 tosses. Given the method of assigning M -values described in the previous section, for appropriate values of q one can expand to K_3 by adding a hypothesis of the form "the frequency of heads is between $500 - k$ and $500 + k$."

My proposal is that the subjunctive "if the coin a were tossed 1,000 times, a would land heads approximately 500 times" is supported by the chance assumption when the subjunctive appraisal involves an appraisal with respect to serious possibility relative to a corpus obtained by first contracting with minimal loss of information, then expanding by adding the antecedent of the counterfactual conditional, and finally making an inferential expansion.

The prima facie conflict between the first and the second counterfactual appraisal is due to the fact that in the second appraisal the corpus used for appraisal with respect to serious possibility involves the first two transformations of the initial corpus but not the third.

Which kind of counterfactual (or subjunctive) evaluation is intended can only be gleaned from context.

The chief point is that we should not be tempted to analyze "has a .5 chance of heads on a toss and a .5 chance of tails" in terms of the relative frequencies with which heads (tails) would occur on long sequences of repeated tosses. Even if we incorporate the stipulation that the repeated tosses be stochastically independent and the chances constant (which already introduces chance predicates) in the counterfactual antecedent, the evaluations of modality would have to be relative to a corpus like K_3 . To obtain K_3 , we have to invoke a principle of direct inference relative to K_2 to obtain the credal state used in making an inferential expansion.

These last moves already presuppose understanding of the concept of chance and, indeed, are essential to that understanding. Any counterfactual analysis of chance can only be understood by already having a grasp of direct inference and inductive expansion.

Thus, we return to the main theme: The principle of direct inference associated with a simple chance predicate is critical to the understanding of that predicate and of the connections between knowledge of chances and judgments of test behavior. It is because of this centrality that we are entitled to consider such principles of direct inference to be principles of inductive logic.

**12.13
Composite
Hypotheses:
The Finite
Case**

Attention has been devoted until now to direct inference from knowledge of simple statistical or chance hypotheses to judgments about test behavior.

Suppose, however, that X knows the following information:

- (1) Setup a is a $CD(\Omega/s)$.
- (2) For some chance distribution F_θ where $\theta \in \Theta$, a is a $C(\Omega/S/F_\theta)$.
- (3) At time t a trial of kind S & T occurs to a .
- (4) The information that the trial is of kind T is stochastically irrelevant.

Let e_α assert that a response of kind R_α occurs where $\alpha \in \mathcal{F}(\Omega)$. For any $\theta \in \Theta$, h_θ asserts that a is a $C(\Omega/S/F_\theta)$. h_θ asserts that a is $C(\Omega/S/F_\theta)$ for some $\theta \in \Theta$. h_θ is the assertion (2).

The problem we are now concerned with is to determine what implications the principle of direct inference has for permissible values of $Q(e_\alpha) = Q(e_\alpha; h_\theta)$.

Partial headway may be made by noting that according to credal coherence every permissible Q -function must satisfy

$$(5) \quad Q(e_\alpha; h_\theta) = \sum_{\theta \in \Theta} Q(e_\alpha; h_\theta)Q(h_\theta; h_\theta) \\ = \sum_{\theta \in \Theta} Q(e_\alpha; h_\theta)Q(h_\theta).$$

Borrowing De Finetti's terminology and modifying it slightly,⁴ I shall call this condition *finite conglomerability*. It derives from the condition of finite additivity of Q -functions imposed in section 4.2.

Confirmational conditionalization and direct inference imply that $Q(e_\alpha; h_\theta) = F_\theta(\alpha)$ for every $\theta \in \Theta$. Direct inference requires that e_α be assigned $F_\theta(\alpha)$ as its unconditional Q -value relative to the expansion of X 's corpus obtained by adding h_θ .

Confirmational conditionalization then determines the result thus cited.

Consequently, from (5) and the principle of direct inference from knowledge of simple chance hypotheses and confirmational conditionalization we obtain

$$(6) \quad Q(e_\alpha) = Q(e_\alpha; h_\theta) = \sum_{\theta \in \Theta} F_\theta(\alpha)Q(h_\theta).$$

Thus, the set of permissible Q -values for the hypothesis e_α depends on the permissible Q -distributions over the simple statistical hypotheses h_θ that are disjuncts in the composite statistical hypothesis h_θ .

To illustrate, suppose that X knows that $h_{.4}$, which asserts that the chance of coin a landing heads on a toss is .4, or $h_{.6}$, asserting that the chance is .6, is true. X also knows that coin a is to be tossed at t . Finally, stochastic irrelevance obtains.

e_H asserts that coin a lands heads on the toss.

Direct inference and confirmational conditionalization imply that $Q(e_H; h_{.4}) = .4$ and $Q(e_H; h_{.6}) = .6$ for every permissible Q -function in the credal state. Credal coherence and this result implies that

$$Q(e_H) = .4Q(h_{.4}) + .6Q(h_{.6})$$

and that

$$Q(h_{.4}) = 1 - Q(h_{.6}).$$

Thus, the set of permissible Q -distributions over e_H and $\sim e_H$ is determined by the set of permissible Q -distributions over $h_{.4}$ and $h_{.6}$.

Credal convexity guarantees that the set of Q -distributions over $h_{.4}$ and $h_{.6}$ is a convex subset of all probability distributions over these alternatives. It follows that the permissible Q -values for e_H constitute a subinterval of the interval from .4 to .6.

These are the strongest results obtainable in this example relying on principles of objectivist inductive logic alone. Objectivist inductive logic (section 4.5) claims that the principles of inductive logic are exhausted by credal coherence and direct inference. The only way more definite results could be derived from these principles is through invoking direct inference from knowledge of chances concerning how coin a was obtained.

Thus, if coin a were drawn at random from an urn containing .4 coins and .6 coins in some definite proportions, direct inference could mandate assigning $Q(h_{.4})$ a Q -value equal to the known proportion of .4 coins in the urn; for, in this context, to select a coin at random is to conduct a kind of trial on the urn which yields .4 coins with a chance equal to their proportion in the urn.

Knowledge of this sort will not, however, always be available and objectivist inductive logic will have nothing left to tell us.

Observe, however, that if X is a necessitarian who thinks one should suspend judgment between Q -functions not forbidden by inductive logic and is, in addition, an objectivist, he will be obliged in our example to assign $h_{.4}$ and $h_{.6}$ all Q -values $(x, 1 - x)$, where $0 \leq x \leq 1$. Consequently, the interval of permissible Q -values for e_H will range from .4 to .6.

Formula (6) shows how to generalize these observations to cases of direct inference from composite statistical hypotheses where the set Θ of simple statistical alternatives is finite.

12.14 Composite Hypotheses: The Countably Infinite Case

In section 12.13, we assumed Θ to be finite. Consider now a situation where Θ is countably infinite.

In this case, condition (5) is equivalent to countable additivity and implies the condition De Finetti called conglomerability—namely, that if the conditional probability of a proposition on each of a countably infinite range of exclusive and exhaustive alternatives falls in a given interval, the unconditional probability also falls in that interval. In section 5.11, the question of violating countable additivity was discussed. I registered agreement with De Finetti's view that we should not be obliged to obey it as a matter of credal rationality or inductive logic.

This circumstance renders the investigation of direct inference from composite hypotheses more complex in the countably infinite case than it was in the finite case. The technical details are summarized in this section. Following the approach of section 5.11, Q -functions in credal states will be represented by σ -finite measures. We shall assume that $m_K(h_\theta) = b_\theta > 0$ and finite for $\theta \in \Theta$.

By direct inference, $Q(e_\alpha; h_\theta) = F_\theta(\alpha) = p_\theta$.

Case (i): $\sum_{\theta=1}^{\infty} b_\theta$ is finite. Hence $\sum_{\theta=1}^{\infty} Q(h_\theta) = 1$. Countable additivity obtains so that conglomerability is satisfied and formula (5) of the previous section can be used. Then

$$Q(e_\alpha) = \frac{\sum_{\theta=1}^{\infty} p_\theta b_\theta}{\sum_{\theta=1}^{\infty} b_\theta}$$

Thus, case (i) can be handled in substantially the same manner as the finite case.

Case (ii): $\sum_{\theta=1}^{\infty} b_\theta = \infty$; $m_K(e_\alpha) = \sum_{\theta=1}^{\infty} b_\theta p_\theta$ is finite. In this case, $m_K(\sim e_\alpha) = \infty$. Hence, $Q(e_\alpha) = 0$.

It can be shown that the glb of the p_θ 's is 0. If it were positive and equal to the value k , then $\sum_{\theta=1}^{\infty} b_\theta k < m_K(e_\alpha)$. But this is impossible since $\sum_{\theta=1}^{\infty} b_\theta = \infty$ and $m_K(e_\alpha)$ is finite. Hence, the glb is equal to 0.

Thus, in case (ii) not only do direct inference, credal coherence, and confirmational conditionalization yield a definite value for $Q(e)$ —namely 0—from the Q -distribution of the h_θ 's; but that value is one of the bounds of the range of values for the p_θ 's.

When $m_K(e_\alpha)$ is infinite while $m_K(\sim e_\alpha)$ is finite, then $Q(e_\alpha) = 1$. In this case, like the one just considered, direct inference from knowledge of composite chance assumptions assigns e_α a Q -value which is one of the bounds of the range of values for the p_θ 's.

To this extent, therefore, case (ii) situations resemble those of case (i). It is now time to turn to cases which deviate more radically from the finite case.

Case (iii): $\sum_{\theta=1}^{\infty} b_\theta = m_K(e_\alpha) = m_K(\sim e_\alpha) = \infty$.

In section 5.11, it was noted that derivation of values for $Q(e_\alpha)$ and $Q(\sim e_\alpha)$ in this case requires appeal to a new σ -finite measure n_K defined over a σ -subalgebra of the original σ -algebra.

If matters are left in this state, we shall not be in a position to determine the value for $Q(e_\alpha)$ from the Q -distribution or m_K -distribution for the h_θ 's, coherence, conditionalization, and direct inference. The function n_K will have to be invoked as well.

The important point is that $Q(e_\alpha)$ might as a result be as-

signed a Q -value that falls outside the interval within which p_θ ranges.

To acknowledge this is to do no more than to admit that conglomerability may be violated once countable additivity is abandoned.

However, there are some special variants of case (iii) where we may derive values for $Q(e_\alpha)$ falling within the range of values for p_θ . It is important that these cases be identified.

It is entertainable that infinitely many h_θ 's specify chance distributions over the sample space which assign the same chance to R_α . Let that chance be q_ϕ . Let g_ϕ be the composite chance hypothesis asserting that exactly one of these h_θ 's is true.

The example of the die discussed in section 12.7 can be used to support the view that the background commentaries for chance predicates guarantee that if we consider the sample space consisting of the events R_α and $\sim R_\alpha$, the simple chance hypothesis f_ϕ asserting that the chance of R_α is equal to q_ϕ is true if any one of the h_θ 's entailing g_ϕ is true. Hence, f_ϕ is true if g_ϕ is true. But direct inference and confirmational conditionalization guarantee that $Q(e_\alpha; f_\phi) = q_\phi$. To secure consistency, we must have $Q(e_\alpha; g_\phi) = q_\phi$.

Case (iiia): Let the h_θ 's specify only finitely many distinct values q_ϕ as chances for R_α . By what has just been said, the h_θ 's can be grouped into finitely many distinct categories corresponding to finitely many hypotheses g_ϕ such that $Q(e_\alpha; g_\phi) = q_\phi$.

Suppose, in the first instance, there is exactly one such value q_ϕ . Since $\sum_{\theta=1}^{\infty} b_\theta = \infty$, then $m_K(e_\alpha) = \sum_{\theta=1}^{\infty} b_\theta q_\phi = q_\phi \infty = \infty$. Similarly, $m_K(\sim e_\alpha) = \infty$. Thus we have a bona fide case (iii) predicament.

Direct inference, however, mandates that $Q(e_\alpha; g_\phi) = q_\phi$. Hence, there must be a suitable σ -finite measure n_K such that $n_K(e_\alpha \& g_\phi)/n_K(g_\phi) = q_\phi$.

A variant on this situation emerges when there is more than one value q_ϕ , although at most finitely many, and where only one q_ϕ contains infinitely many disjuncts. It can then be shown in case (iii) that only the $m_K(g_\phi)$ with infinitely many disjuncts has infinite value. Hence, that hypothesis obtains Q -value of 1 and $Q(e_\alpha) = Q(e_\alpha; g_\phi) = q_\phi$ for that particular value of ϕ .

Case (iiib): This is like case (iiia) except that more than one but finitely many of the g_ϕ 's are disjunctions of infinitely many h_θ 's. This case parallels in a straightforward manner the finite case of the previous section.

Case (iiic): The h_θ 's specify infinitely many distinct p_θ 's and, hence, q_ϕ 's. There are infinitely many g_ϕ 's each of which is a disjunction of finitely many h_θ 's.

Case (iiid): As (iiic) except that one g_ϕ is a disjunction of infinitely many h_θ 's.

Case (iiie): As (iiic) except that finitely many g_ϕ 's are disjunctions of infinitely many h_θ 's.

Case (iiif): As (iiic) except that infinitely many g_ϕ 's are disjunctions of infinitely many h_θ 's.

Cases (iiid) and (iiie) reduce to cases (iiia) and (iiib), respectively. Case (iiif) reduces to distinct cases parallel to cases (i), (ii), and (iii), only now with n_K -functions replacing the m_K -functions.

We are left, therefore, with case (iiic). Let $m_K(g_\phi) = b_\phi$, which is positive and finite. So is $m_K(e_\alpha \& g_\phi)$. Moreover, $m_K(e_\alpha \& g_\phi)/m_K(g_\phi) = q_\phi$ (or a positive multiple thereof) by direct inference.

On the other hand, $m_K(e_\alpha) = m_K(\sim e_\alpha) = \infty$. The σ -finite measure m_K determines, therefore, the values of $Q(e_\alpha; g_\phi)$ and $Q(g_\phi)$ for each g_ϕ . The first value is q_ϕ and the second is 0. Hence, $\sum_{\phi=1}^{\infty} Q(e_\alpha; g_\phi)Q(g_\phi) = 0$. Yet, $Q(e_\alpha)$ is left entirely undetermined. Its value depends on how the function n_K is chosen.

Thus, in case (iiic) it is possible to pick a value of $Q(e_\alpha)$ outside the interval within which values of p_θ or q_ϕ fall without violating direct inference.

Case (iiic) situations rarely loom large, and throughout the remaining discussion we shall have no occasion to consider them. Yet, it is important to understand that in principle circumstances may arise where a numerically precise credal state over rival simple chance hypotheses fails to yield a definite Q -value for e_α through the offices of direct inference and confirmational conditionalization.

Such cases may be bizarre. That is my excuse for ignoring

them in what follows. However, someone might identify important or interesting situations meeting the specifications of case (iiic) in future applications. They may prove less bizarre than I now think.

**12.15
Composite
Hypotheses:
The Continu-
ous Case**

Let us assume that the simple chance hypotheses in h_θ constitute a noncountably infinite set indexed by n -tuples θ of real numbers belonging to some finite region of an n -dimensional space. If, as before, p_θ is the chance of obtaining result R_α on a trial of kind S according to h_θ , then $Q(e_\alpha)$ will equal $\int p_\theta f(\theta) d\theta$ provided we have a suitable continuous density function $f(\theta)$ over the points in the region Θ .

Of course, I am supposing that in this situation we do have a probability measure which obeys countable additivity at least for the set of Lebesgue measurable sets in the region.

This requirement of countable additivity, however, is abandoned for some situations where Θ is the entire space. Thus, if θ ranges from $-\infty$ to $+\infty$, and if $f(\theta)$ is constant, then $f(\theta)$ must equal 0; and this leads to a violation of countable additivity.

The technical complexities which emerge parallel those discussed in the previous section, and will not be repeated here.

**12.16
On a Pseudo
Paradox**

Violations of countable additivity occur in those cases where X wishes to adopt as his credal state for countably many alternative hypotheses one which assigns each alternative equal Q -value. Thus, if X considers each value of the mean of a normally distributed random variable to be seriously possible and partitions the space of possible values into intervals of equal finite length, he might wish to regard the hypothesis that the true value falls in one such interval as equally probable as the hypothesis that the true value falls in another such interval.

Objections have been raised to the use of such distributions on the grounds that they breed paradoxical results and, in recent years, a small industry has developed in the production of such paradoxes. M. Stone is one of the leading contributors and it may be useful at this point to consider a relatively elementary example propounded by Stone himself.⁵ My aim is to show that the paradox is obtained by appealing to countable additivity at a critical stage in the argument. Since countable additivity is violated at the outset by adopting a uniform

distribution over countably many alternatives, it is not surprising that a contradiction can be generated. To use this fact as an objection to using so-called improper distributions, however, is to beg the question under dispute.

De Finetti went over the same ground many years ago with more elementary and clearer examples.⁶ I go through the exercise here simply to point out that his entirely cogent position is not refuted by hiding behind the skirts of complexity.

Consider a plane marked off in a grid whose intersections are the points (x, y) , where x and y are taken to be integers. A point (w, z) is adjacent to (x, y) if and only if either $w = x \pm 1$ and $z = y$ or $w = x$ and $z = y \pm 1$. There are, therefore, four points adjacent to any given point.

A path θ is a finite sequence of points such that the points preceding and succeeding a given point in the sequence are adjacent to it—if such points exist (i.e., if the point in question is not an endpoint).

For a sequence of n points, where $n \geq 1$, the length of the sequence is $n - 1$; so the sequence can have any nonnegative finite value. Every sequence has a first point and a last point. There are a countable infinity of distinct paths.

Let an endpoint of a path θ be identified by (a, b) . On a trial, one of four outcomes can occur with a chance of $1/4$. If the outcome is 1, 2, or 3, the path is extended to one of three adjacent points not already on θ in accordance with a definite procedure for deciding which of the three extensions to make on the basis of the result of the trial.

If the outcome is 4, the path is shortened by moving back on the path to the point adjacent to (a, b) already on the path.

In the special class of cases where the length of a path is 0 (there are infinitely many of these), the path is lengthened in each of the four possible ways depending on which of the four outcomes of the trial takes place.

Thus, for every value of θ , the chance of lengthening θ on a trial of this sort is $3/4$ for all θ 's of positive length and 1 for all θ 's of 0 length.

Suppose X knows that a trial is to be conducted on some specific occasion, but neither the value of θ nor the outcome of the trial is known.

However, X knows the hypothesis e which states which path results from conducting the experiment.

Let h assert that a lengthening has taken place and $\sim h$ assert that a shortening has taken place. Let g assert that the length of the true θ is positive and $\sim g$ assert that it is 0.

By direct inference, $Q(h; g) = 3/4$ and $Q(h; \sim g) = 1$. Depending on the permissible Q -values for g and $\sim g$, $Q(h)$ takes a value somewhere from $3/4$ to 1.

Recall that X knows e . Hence there are four hypotheses as to the true value of θ consistent with what he knows.

One of these implies that a lengthening has taken place and the other three imply that a shortening has taken place except when the path described by e has 0 length. In that case, X knows for sure that a shortening has taken place.

By direct inference, it can be shown that for each of the four values of θ compatible with the information specified by e , $Q(e; \theta) = 1/4$ according to the credal state prior to finding out that e is true.

Suppose that X 's credal state prior to finding out the truth of e assigns each hypothesis as to the true value of θ equal Q -value and, hence, 0 Q -value. That is to say, we have a σ -finite measure m_κ such that $m_\kappa(\theta) = c$ for all θ . Hence, $m_\kappa(e \& \theta)$ must equal $c/4$ for values of θ consistent with e and otherwise 0. Hence, $M_\kappa(e) = c$. Thus, for all e 's, $Q(\theta; e) = 1/4$ if θ is consistent with e and 0 otherwise.

Recall, however, that for e specifying a path of positive length, there is exactly one θ for which that path is a lengthening. Hence, $Q(h; e) = 1/4$. If e describes a path of 0 length, $Q(h; e) = 0$.

Thus far, there is nothing to dispute about the reasoning.

Consider the hypothesis f asserting that one of the countable infinity of hypotheses of type e specifying a path of positive length is true and let $\sim f$ assert that the path resulting from the trial has a 0 length.

What are the values of $Q(h; f)$ and $Q(h; \sim f)$?

We know that $Q(h; \sim f) = 0$ simply because h is inconsistent with f .

What about $Q(h; f)$?

We cannot say that it equals $1/4$ by direct inference. If, however, we invoke countable additivity), $Q(h; f)$ is equal to the countable sum of all products of the type $Q(h; e)Q(e)$ for all e specifying a positive length and this should lead to $Q(h; f) = 1/4$. Conglomerability will have been established.

If this were the case $Q(h)$ would be confined to values between 0 and $1/4$. That is in contradiction to our previous argument which confined the value to the range between $3/4$ and 1.

Stone calls results of this sort examples of "strong inconsistency."⁷

There is no inconsistency here unless one insists on mandating conglomerability as a general condition on Q -functions. We have, therefore, a choice between giving up the uniform "prior" distribution over the paths θ or abandoning conglomerability and thereby rejecting countable additivity. But anyone prepared to allow uniform distributions is prepared to allow violations of countable additivity and presumably conglomerability as well.

Suppose X is offered 1 dollar for 3 in a gamble on the truth of h . He is to decide whether to accept or reject the gamble after finding out the path which results from the experiment (so that he knows which e is true).

Given that information, X should assign h a degree of credence equal to $1/4$ or 0 depending on whether the length of the path is positive or 0. He should, therefore, reject the gamble as unfavorable to him.

On the other hand, had X been constrained to plan which option to choose depending on the path e he discovers after the experiment is performed but to do so before finding out which e is true, the chance of winning by accepting the bet regardless of what he finds out is at least $3/4$; so that he should be prepared to accept the bet regardless of what happens.

Stone finds this result objectionable. It has driven former users of improper distributions to abandon their practice.⁸

The counterintuitive air surrounding the betting phenomenon derives from an illicit extrapolation from what is objectionable in cases where considerations of infinity are not involved to the infinite case. De Finetti puts the matter well when he writes:

From Bolzano (*Paradoxien des Unendlichen*) to Borel (*Les Paradoxes de l'Infini*) every new property of infinity has been considered paradoxical. Cantor's ideas were for a long time rebutted: logical paradoxes pushed logic into a quandary. No wonder if our case is no exception.⁹

13.1
Direct
Inference
and Bayes'
Theorem

Chapter 12 focused on a family of principles of direct inference designed to regulate adoption of hypotheses about test behavior on the basis of knowledge of chances. Section 12.11 reviewed some procedures for statistical prediction based on criteria for deliberate inductive expansion. In such expansion, knowledge of chances may be used to justify predictions about relative frequencies in large sequences of trials.

To have a grasp of the role of chance in inquiry and deliberation also requires an account of how knowledge of chances may be acquired on the basis of information about test behavior. Given a corpus K_e containing information about test behavior and an ultimate partition U of rival statistical hypotheses, which elements of U may be eliminated via deliberate inductive expansion on the basis of the evidence in K_e and the investigator's demands for information?

The view of inferential expansion proposed earlier requires appeal to X 's credal state for elements of U relative to the corpus K_e . Of course, that credal state is determined by K_e and X 's confirmational commitment. Consequently, the legitimacy of an inferential expansion in which X adds information about chances to his corpus depends on the corpus K_e and X 's confirmational commitment. This dependence points to the desirability of exploring how knowledge of test behavior contained in K_e determines through the confirmational commitment what X 's credal judgments for rival statistical hypotheses should be.

Let K_e be the deductive closure of K and information e concerning the outcome of some experiment and consider X 's credal state $C(K)$ relative to K . Let $h \in U$ and, for the sake of simplicity, consider a Q -function in $C(K)$ for which $Q(e) > 0$. The multiplication theorem embedded in credal coherence implies that

$$Q(h; e) = Q(e; h) \frac{Q(h)}{Q(e)},$$

where, following the jargon of section 4.3, $Q(h; e)$ is the posterior Q -value for h relative to e , $Q(h)$ is the prior Q -value for h , and $Q(e; h)$ is the likelihood of h on e according to that particular Q -function in $C(K)$.

According to Bayes' theorem (section 4.3), credal coherence requires that $Q(h; e)$ be determinable by the prior Q -value for h , its likelihood on e , and the prior Q -values and likelihoods on e of all other elements of U . This last result derives from the fact that $Q(e) = \sum_{h' \in U} Q(e; h')Q(h')$.

$Q(h; e)$ is a permissible posterior Q -distribution over the elements of U according to the credal state $C(K)$. For any such permissible posterior distribution, let $Q'(h) = Q(h; e)$ where h is a member of U . Confirmational conditionalization requires that $Q'(h)$ define a permissible unconditional distribution over the elements of U according to $C(K_e)$. Thus, posterior distributions determined via Bayes' theorem for the credal state $C(K)$ become unconditional distributions in $C(K_e)$ through the application of confirmational conditionalization.

Of course, if X begins with confirmational commitment C and corpus K , conducts an experiment, and adds e to his corpus to form K_e , he might also revise his confirmational commitment by shifting from C to C' . His new credal state $C'(K_e)$ will not be the conditionalization of $C(K)$. But it will be the conditionalization of $C'(K)$. It may be the case that the warrant for a shift from C to C' derives from a comparison of the properties of $C(K)$ and $C'(K)$, and, in particular, a comparison of the permissible prior Q -distributions for elements of U and permissible likelihood functions on the possible outcomes of experiments. Thus, even if we take into account that temporal credal conditionalization may break down in shifting from K to K_e , it is important to understand the ramifications of Bayes' theorem.

The application of Bayes' theorem depends on determination of two factors: prior credal states for elements of U and, for each hypothesis describing an outcome of experimentation, the set of permissible likelihood functions relative to that hypothesis.

The principle of direct inference can play an important role in the determination of both of these factors.

To illustrate, consider the following artificially constructed example:

X knows that urn b contains some coins biased in favor of heads with a .9 chance of landing heads on a toss of kind S and some coins biased in favor of tails with a .1 chance of landing heads on a toss of kind S . The percentage of coins of the first type is $100x\%$ and of the second type is $100(1 - x)\%$.

X selects a coin from urn b at random—i.e., by a method known to select coins biased in favor of heads with a chance equal to the relative frequency x of such coins in urn b . Let that coin be a .

X is now concerned to find out whether coin a is biased in favor of heads ($h_{.9}$) or tails ($h_{.1}$).

X will toss coin a in manner S exactly once and observe whether the coin lands heads (e_H) or not.

A toss of kind S is conducted by selecting a tossing device at random from urn c of such devices. Some devices in urn c are of type A and others are of type B . Some are red and others are blue. The chance of heads on a toss of a coin biased in favor of heads with a red type A device or a blue type B device is .98. With a blue type A device or a red type A device, it is .1. The chance of heads on a toss biased in favor of tails with a red type A device or a blue type B device is .1, and with a blue type A device or a red type B device is .98.

The percentage of devices of type A that are red or type B that are blue must be 90.91%. Otherwise, the chance of heads on a trial of kind S would not be .9 for coins biased in favor of heads.

When X conducts a trial of kind S on coin a , he will find out the color of the tossing device but not whether it is of type A or of type B . Thus, he will know not only that the trial is of kind S but also kind T_R or T_B depending on whether the tossing device used is red or blue.

If, however, the percentage of type A devices in urn c among the red ones is equal to the percentage of red type A or blue type B devices—i.e., 90.91%—the information concerning the color of the device will be stochastically irrelevant. Suppose X does know this.

The principle of direct inference obligates X to assign $h_{.9}$ the

precise value x as uniquely permissible relative to X 's knowledge prior to finding out whether e_H is true: $Q(h_{.9}) = x$. The same principle obligates X to adopt a credal state where, for all permissible Q -values, $Q(e_H; h_{.9}) = .9$ and $Q(e_H; h_{.1}) = 1$.

According to Bayes' theorem,

$$Q(h_{.9}; e_H) = \frac{.9x}{.9x + .1(1 - x)}.$$

When X finds out that e_H is true and adds this information to form a new corpus K_H , the principles of inductive logic and confirmational conditionalization obligate him to assign $h_{.9}$ the value $.9x/[.9x + .1(1 - x)]$ as the uniquely permissible unconditional Q -value relative to K_H . This is so regardless of other changes he may or may not make in his confirmational commitment. In this sense, his knowledge and the principles of inductive logic dictate what his posterior credal state should be because they determine what his prior credal state for the rival hypotheses should be and what his likelihood function on the data should be.

The inductive logic used is an objectivist inductive logic in the sense of section 4.5. The sole principles of inductive logic invoked are credal coherence and direct inference.

Of course, the results obtained depend on more than inductive logic. X knows (accepts as evidence in his initial corpus K) a substantial amount of information about chances. He knows that coin a was selected at random from urn b and also knows the chance of obtaining a coin biased in favor of heads on such a trial. He knows that a is tossed in manner S with a red tossing device. He also knows that the chance of heads on a toss of a in manner S with a red tossing device is equal to the chance of heads on a toss of coin a in manner S regardless of whether a is biased in favor of heads or tails. That is to say, he knows the information concerning the color of the tossing device to be stochastically irrelevant.

With such information, he is in a position to invoke direct inference to determine prior credal probabilities and likelihoods.

There are two ways in which deprivation of background information about chances could frustrate the application of Bayes' theorem for the purpose of deriving posterior credal

distributions for rival statistical hypotheses on the basis of data:

(a) X 's corpus may not warrant deriving a numerically precise prior credal distribution for the rival statistical hypotheses.

(b) X 's corpus may not warrant deriving a unique likelihood function for the rival hypotheses on the data for all possible outcomes of experimentation.

Breakdown of one sort or another does not by itself preclude adopting numerically precise priors and likelihoods. However, such priors and likelihoods would not be mandated by objective inductive logic and the available knowledge alone.

Critics of the use of Bayes' theorem to derive posterior probabilities from priors, such as R. A. Fisher, have not claimed that such uses are always illegitimate. Indeed, Fisher himself has cited instructive illustrations where such application is legitimate.

Fisher considers a situation where a black mouse which is born to black heterozygous parents is mated with a brown mouse (which must be homozygous, because the color brown is recessive). The chance of a black offspring of two heterozygous parents being homozygous is $1/3$. Hence, if one wished to test to find out whether the black mouse is homozygous by investigating the color of its offspring with the brown mouse, direct inference provides a prior credal probability of $1/3$ for the hypothesis that the mouse is homozygous black. Fisher considers how the likelihood function for the two hypotheses (the mouse is homozygous and the mouse is heterozygous) can be determined on the basis of the observation of seven offspring by appeal to direct inference. He then explains how Bayes' theorem can be used to compute a posterior distribution for the two hypotheses from the prior credal distribution and the likelihood function. Fisher then writes:

If, therefore, the experimenter knows that the animal under test is the offspring of two heterozygotes, as would be the case if both parents were known to be black, and a parent of each were known to be brown, or if, both being black, the parents were known to have produced at least one brown offspring, cogent knowledge *a priori* would have been available, and the method of Bayes could properly be applied. But,

if knowledge of the origin of the mouse tested were lacking, no experimenter would feel he had warrant for arguing as if he knew that of which in fact he was ignorant, and for the lack of data Bayes' method of reasoning would be inapplicable to his problem.¹

Fisher's remarks are instructive in several ways.

He sees no objection to utilizing Bayes' theorem in deriving posterior credal probabilities provided that all the information required for deriving priors (and likelihoods) via direct inference is available.

Furthermore, when such information is not available, he objects to the adoption of numerically precise priors without derivation from knowledge of chances via direct inference. The applications of Bayes' theorem Fisher objects to (in the passage cited) are those designed to obtain numerically precise posterior probabilities from numerically precise prior probabilities without any warrant for adopting precise priors grounded in knowledge of chances.

Fisher is quite clearly committed in his writings to a distinction between objective statistical probability or chance, on the one hand, and judgments of what I have called credal probability on the other. His characterizations of both ideas are sketchy and, hence, I cannot claim that his way of drawing the distinction coincides in all details with mine. But there is sufficient clarity in what he says to justify interpreting him as committed to some principle of direct inference and to restricting the justified assignment of numerically precise credal probabilities to hypotheses only when such justification is grounded on knowledge of chances and ignorance about features of the kinds of trials being conducted that would prevent such judgments of credal probability.²

J. Neyman appears to take a more restricted view, allowing only a notion of objective or statistical probability to be meaningful. Yet, Neyman would, no doubt, acknowledge that if agent X has knowledge of such objective probability or chances specifying that the chance of heads on a toss is p and knew that the coin is to be tossed, the agent should adopt p as a "fair betting quotient" for bets on the outcome of the toss. I prefer to understand his rejection of the meaningfulness of probability assignments to hypotheses as a rejection of the legitimacy of assigning numerically precise credal probabilities

to hypotheses unless these are derivable via direct inference from knowledge of chances. When so understood, his stand is close to Fisher's on this particular matter. Bayes' theorem may be used to derive numerically precise posterior credal probabilities from numerically precise priors when the priors can be derived via direct inference from knowledge of chances. Otherwise Bayes' theorem cannot legitimately be used in this way.³ I think E. S. Pearson's outlook on this issue is substantially the same.⁴

In making these observations, I do not intend to mask the important differences between the views of Fisher, on the one hand, and Neyman and Pearson on the other. I am interested in emphasizing a shared outlook toward the limitations in applying Bayes' theorem to the derivation of numerically precise posteriors on the basis of experimental data from precise prior credal distributions over rival statistical hypotheses.

Both the concession that Bayes' theorem can be used in this way provided the prior distribution (and the likelihood function) can be derived via direct inference from knowledge of chances and that it cannot be so employed when such knowledge is lacking are basic doctrine for many students of statistical theory.

These two claims are implied by the view I called "objectivist necessitarianism" in section 4.5. This position insists that an objectivist inductive logic consisting of credal coherence and direct inference is a complete inductive logic. It also insists that relative to corpus K no Q -function should be counted impermissible provided that it conforms to the dictates of inductive logic.

Consequently, when knowledge of chances fails to justify adoption of a numerically precise prior via direct inference, objectivist necessitarians prohibit the adoption of such a prior on the grounds that there is no justification for such a choice and the choice would, therefore, be arbitrary. On the other hand, if the requisite knowledge of chances is available, Bayes' theorem and confirmational conditionalization, together with the fact that necessitarians are obliged to remain faithful to their confirmational commitment, entail the derivation of a numerically precise posterior.

Neither Fisher nor Neyman and Pearson endorse every jot and tittle of objectivist necessitarian doctrine.

Thus, in cases where direct inference fails to yield a numerically precise prior from the available knowledge of chances, none of these authors explicitly say that all prior Q -distributions obeying credal coherence are permissible. They can just as readily be construed as denying the permissibility of any Q -distribution in such cases. The former interpretation is the charitable one, as section 5.8 explains; but the historical record does not always support charity.

These authors undoubtedly endorse some version of the principle of direct inference. Whether they accept the one I favor and have built into objectivist necessitarianism is another, and an open, question. It is also unclear whether they would permit violations of confirmational conditionalization.

Further speculation on the views of these authors will be postponed until later chapters. My concern now is to note that objectivist necessitarianism does capture some widely held views concerning derivation of posteriors from priors via Bayes' theorem. Moreover, objectivist necessitarianism would not be wholeheartedly endorsed by anyone, because it has some embarrassing consequences. Outlooks such as Fisher's or the family of views derivative from the Neyman-Pearson approach may be understood as attempts to avoid the embarrassing consequences while severely limiting the applicability of Bayes' theorem in the manner indicated above.⁵

Perhaps no one has ever been a strict objectivist necessitarian (just as no one, in his heart of hearts, is a strict Bayesian). But the objectivist necessitarianism represents a tendency which is quite attractive. In the first place, the only principles of inductive logic which it endorses are principles for which some sort of rationale can be found in the applications of probability in deliberation and inquiry. In the second place, necessitarianism avoids the excesses of personalism by forbidding rational men to adopt one numerically precise credal state rather than another when there is no warrant for doing so.

As I stated in chapter 4, my chief reservation with objectivist necessitarianism is its presupposition that the only sort of warrant for ruling out a Q -function relative to K is failure to satisfy the requirements of inductive logic. Later in this chapter, I shall explain some further considerations which may, on some occasions, warrant ruling out other Q -functions.

Even so, objectivist necessitarianism would remain a compelling doctrine and my reservations mere carping criticism were it seriously applicable in scientific inquiry and practical deliberation.

It is not applicable or, at the very least, there is a *prima facie* case that is not. In the following section, I shall explain the basis for this claim. The rest of the chapter will then be given over to an outline of those fragments of a positive account of the revision of credal states and confirmational commitments that I can offer at present.

The chapters that follow will then consider a few responses to the predicament of objectivist necessitarianism alternative to my own—including responses that to some degree are reconstructions of the views of Fisher on the one hand, and Neyman and Pearson on the other.

13.2 Objectivist Necessitarianism and the Relevance of Data

Consider the example of the coin a in section 13.1. Let us suppose that X knows that the color of the tossing device used on a trial of kind S is stochastically irrelevant, so that direct inference can be used to determine a uniquely permissible likelihood function. Take $Q(e_H; h_{.9}) = .9$ and $Q(e_H; h_{.1}) = .1$.

However, unlike the situation in the original example, X does not know how coin a was obtained and, hence, lacks any knowledge on the basis of which to determine a numerically precise prior Q -value $Q(h_{.9}) = x$ via direct inference.

According to an objectivist necessitarian, every value for x from 0 to 1 is permissible relative to K . Bayes' theorem implies that $C(K)$ contains Q -functions for which $Q(h_{.9}; e_H) = y$ for $0 \leq y \leq 1$. Confirmational conditionalization implies that $C(K_H)$ allows all values for z where $0 \leq z \leq 1$ to be assigned as unconditional Q -values for $h_{.9}$. Since the confirmational commitment is CIL (section 4.5), there is no change in confirmational commitment as X expands from K to K_H as a result of observing the outcome of his experiment.

The upshot is that the outcome of the experiment has no effect on X 's unconditional credal state for the two hypotheses $h_{.9}$ and $h_{.1}$.

The same result would emerge if the experimenter were to toss the coin a in manner S a large number of times, where it is known that repetitions are stochastically independent and that chances are constant on each trial.

Nor is the result altered if the set U of rival simple statistical hypotheses h_θ specifying chance distributions of type F_θ over sets in $\mathcal{F}(\Omega)$ contains more than two and even infinitely many elements, or if the sample space contains more than two and even infinitely many points.

Unless the data of experimentation entails the falsity of some element of U , such data will be confirmationally irrelevant for the elements of U relative to K in the sense of section 10.5. To be sure, that data will remain confirmationally relevant in the strong sense of section 10.6; but the conditions under which strong confirmational relevance entails confirmational relevance do not obtain when objectivist necessitarianism is endorsed.

Prima facie, therefore, objectivist necessitarianism implies the uselessness of data concerning test behavior for modifying credal states for hypotheses about chances except under special circumstances when prior knowledge justifies relatively determinate prior credal states via direct inference. This alleged implication, if sound, constitutes a decisive objection to objectivist necessitarianism. Those who are sympathetic to its intent must find some way to modify or repair the damage.

One possibility is to concede that the data are confirmationally irrelevant, but to deny that they are therefore useless. I think that some such view as this is favored by Neyman and Pearson and their followers. I shall consider views such as this in chapter 17.

Another response is to remain necessitarian, but to strengthen inductive logic by adding new principles. However, unlike the classical Bayesians from Bayes himself to H. Jeffreys and R. Carnap, someone sympathetic to objectivist necessitarianism might restrict such additional principles to those that might be thought to articulate incorrigible features of the use of concepts of chance in inquiry and deliberation. This view has been advanced by I. Hacking. I shall consider it in chapter 15.

An alternative strategy is to modify or reject confirmational conditionalization. H. E. Kyburg has proposed doing so through modifying the principle of direct inference. I suspect his outlook resembles Fisher's view more than other reconstructions do. A. P. Dempster has also recommended modi-

fying confirmational conditionalization. Kyburg's and Dempster's approaches will be discussed in chapter 16.

Instead of tampering with objectivist inductive logic, one might abandon necessitarianism in favor of personalism of either the intemperate or the tempered varieties. I shall not reiterate my opposition to such views here.

Finally, one can explore ways and means of exploiting contextual factors additional to the available knowledge for the purpose of adjusting credal states and modifying confirmational commitments. Elements of an objectivist but revisionist view of this sort will be sketched in the remainder of this chapter.

13.3 Likelihood and Irrelevance

The use of Bayes' theorem and confirmational conditionalization to derive posterior credal states conditional on the outcome of experimentation depends not only on the prior credal state for the rival elements of U but on the permissible likelihood functions on the data. In the previous discussion, attention was focused on the problem of identifying the appropriate prior credal state for the purpose of using Bayes' theorem; and this same problem will preoccupy us in most of the remaining portions of the discussion.

It must not be forgotten, however, that the determination of likelihoods also poses problems.

These problems may appear to be less severe even for an objectivist necessitarian when the elements of U are alternative simple statistical hypotheses specifying rival chance distributions over a sample space Ω on trials of a given kind S . The likelihood of the hypothesis h_θ asserting that the chance distribution is F_θ on the datum e_α asserting that result R_α (for $\alpha \in \mathcal{F}(\Omega)$) occurs on the trial of kind S in question may be held to be equal to $F_\theta(\alpha)$, on the basis of direct inference and confirmational conditionalization.

But matters are rarely so simple as this. In real life, X will know more about the kind of trial than that it is of kind S . He will have additional information to the effect that the trial is of kind T . Unless X knows that regardless of which θ is the true one the information that the trial is of kind T is stochastically irrelevant, the principle of direct inference will fail to obligate X to assign likelihoods in the manner just indicated.

Indeed, for an objectivist necessitarian matters are still

worse than this. If X does not know whether the extra information is stochastically irrelevant or not, he will not know what the chance of a result of kind R_α is on a trial which is both an S and a T when h_θ is true. Relative to the initial corpus K , $CIL(K)$ could conceivably allow permissible $Q(e_\alpha; h_\theta)$ to take any value from 0 to 1. This could happen even if the prior credal state for the h_θ 's is numerically determinate. As a consequence, the posterior credal state determined by Bayes' theorem would be indeterminate—indeed, maximally so in some situations, in spite of the numerical precision of the prior. Acquisition of new data would be harmful, rather than helpful or merely irrelevant.

Return to the example of coin a discussed in section 13.1. Suppose now that X knows that coin a was selected from urn b at random and that exactly 50% of the coins in urn b are biased in favor of heads and exactly 50% are biased in favor of tails, so that the priors for $h_{.9}$ and $h_{.1}$ are determined by direct inference. X also knows that the toss of coin a is of type S and that the tossing device used is a red one.

Unlike the previous case, however, X lacks any information as to the chances of heads on a trial of kind S with a red type A or a blue type B device when $h_{.9}$ is true or when $h_{.1}$ is true. He also lacks such information for trials of kind S with a blue type A or red type B tossing device. Finally, he knows nothing of the percentages of red type A , blue type A , red type B or blue type B devices in urn c . All he knows is that the coin a is tossed with the aid of a tossing device selected from urn c and that the device so obtained is red and, in addition, that the chance of heads on a toss with the aid of a device selected at random from urn c is .9 if $h_{.9}$ is true and .1 if $h_{.1}$ is true.

For all X knows, the color of the tossing device may be stochastically irrelevant. But, for all he knows, it may be stochastically relevant. In his ignorance, $Q(e_H; h_{.9})$ can take any value from 0 to 1. And $Q(e_H; h_{.1})$ is equally indeterminate. In particular, the Q -function for which $Q(e_H; h_{.9}) = 1$ and $Q(e_H; h_{.1}) = 0$ and the Q -function for which $Q(e_H; h_{.9}) = 0$ and $Q(e_H; h_{.1}) = 1$ are both permissible.

Bayes' theorem allows, therefore, $Q(h_{.9}; e_H)$ to take any value from 0 to 1 even though the prior is numerically precise.

Needless to say, if X began with background knowledge of

the sort just described, he would see little point in conducting a trial of kind S on coin a to find out whether $h_{.9}$ or $h_{.1}$ is true.

To be sure, if X endorsed a confirmational commitment according to which, conditional on $h_{.9}$, say, being true, the information that the trial is one with a red tossing device is confirmationally irrelevant in the strong sense for deciding whether e_H is true or false relative to the information that the trial is a toss of a of kind S , he could then invoke direct inference to establish that $Q(e_H; h_{.9}) = .9$ and $Q(e_H; h_{.1}) = .1$. But this restriction on X 's confirmational commitment cannot be justified by an appeal to direct inference and coherence alone.

Thus, an objectivist necessitarian cannot invoke such an appeal to confirmational irrelevance to fill the gap which ignorance of stochastic irrelevance leaves; and even those who are not objectivist necessitarians must offer some basis for such a restriction on confirmational commitments.

The coin example used to illustrate the point was an artificial one; but the problem is not artificial.

Suppose X is interested in comparing the effectiveness of two different kinds of manure for the growing of wheat of a certain kind. He divides a number of fields into two plots. He undertakes to make the two plots in any given field as similar in all discernible respects as he is able. On each field, he spreads one type of manure in one plot and the other type in the other. He then plants similar amount of seed in each field and assures cultivation under as uniform conditions as is feasible. At harvest, X then determines the difference in yield for pairs of plots in each field and forms the average of the differences. This data together with data concerning the observed variation in the differences can then be used to compute a value of the so-called Student's t -statistic.

If there is no difference in the efficacy of the two methods of manuring the fields (the "null hypothesis") the chance distribution for values of t on trials of the sort just described is determined. Hence, on experimental data concerning the value of the t -statistic, the likelihood of the null hypothesis is determined—provided that the extra information available to the experimenter concerning the kind of trial is known to be stochastically irrelevant.

Now the experimenter does know, for each field, in which

plot (the "left" or the "right" one) which manure was spread. He may have other such differentiating knowledge as well. Even though the experimenter may have taken all precautions practicable to guarantee that there are no discernible differences between the left and right plots in any given field, he may, nonetheless, not be prepared to take for granted that which plots received which treatments is stochastically irrelevant information. Indeed, in some agricultural experiments, there may be some prior background information indicating that such information about treatments in plots is stochastically relevant without there being clear information as to how it is relevant.

One possible way to confront such difficulties is to engage in preliminary inquiries aimed at identifying such sources of systematic error (e.g., by conducting "uniformity trials" through treating both plots in every field with the same manure or with no manure in order to ascertain peculiarities in the plots which might lead to systematic error.) If such efforts are feasible and successful, one might then hope to correct for such systematic errors when conducting the experiment.

Fisher pointed out, however, that such approaches are often not practicable and even when they are feasible may be too costly.⁵

To adopt a confirmational commitment which secured the confirmational irrelevance of the information as to which treatments were made in which plots in which fields would, under the circumstances, be to adopt a confirmational commitment much stronger than objectivist necessitarians can allow.

Fisher's remedy was to recommend artificial randomization.⁶ That is to say, the choice of which manure is to be spread in which plot is to be based on the outcome of the toss of a coin known to be fair or on some other experiment on a chance setup that selects both plots for a given treatment with equal chance.

Fisher apparently believed that through randomization, information about kinds of trials which could not otherwise be ignored could be ignored. And there is, indeed, a sense in which Fisher is right.

Prior to finding out the results of randomization, the chance distribution for the t -statistic on the null hypothesis is perfectly determinate and if the experimenter could ascertain the value

of that statistic without ever finding out which plots were selected for which treatment as a result of randomization, direct inference would justify a numerically definite likelihood for the null hypothesis (and any simple alternatives to it).

In practice, however, information as to which plots received which treatments is known and, indeed, used when computing the value of the t -statistic. And once it is known, the mere fact that the treatments were assigned to plots at random contributes nothing to establishing the stochastic irrelevance of the information. Without such knowledge, however, direct inference cannot justify definite likelihood for the null hypothesis or any other simple alternative to it.

Thus, it seems to me that the technique of artificial randomization used in comparative experiments (and in other sorts of experimentation such as sample surveys) has very little to recommend it insofar as the rationale for it is the one which Fisher, who is more responsible than anyone for its widespread use in experimental design, offered for it.

To be sure, randomization has other virtues. Like double-blind techniques, systematic arrangements (which are the opposite of random arrangements), and the like, randomization removes the experimenter from the experiment in ways that help prevent unwitting bias or contamination—especially in medical research. But randomization is not the sole means for preventing such contamination, and this cannot account for the often strident insistence that randomization is a *conditio sine qua non* of a good experimental design.

Were Fisher's rationale cogent, such a claim could be made with justification. My own view is that Fisher's argument is not cogent, but that the problem to which he directed attention is an important one. In any experiment, the experimenter will have a considerable amount of extra information about the kind of trials being conducted. On some occasions, his background knowledge will suffice to secure the assumption of stochastic irrelevance, or at least of approximate stochastic irrelevance, so that on data obtained concerning the outcome of the experiment a fairly definite likelihood function can be constructed in conformity with the dictates of objectivist necessitarianism.

On other occasions, X 's background knowledge will not suffice for the purpose. In such cases, either X will have to

look for an improved experimental design, conduct side investigations to ascertain whether the extra information is stochastically relevant or irrelevant and the extent to which it is, or he will have to abandon objectivist necessitarianism. Obviously these are not exclusive alternatives.

In any case, it should be apparent from all of this that in an investigation aimed at finding out which of rival simple statistical hypotheses in U is true on the basis of some experiment or series of experiments, the use of direct inference to obtain a definite likelihood function on data depends on a substantial amount of background knowledge concerning chances.

Such background knowledge includes specification of the sample space for trials of kind S , the set of rival simple statistical hypotheses h_θ specifying chance distributions F_θ over the sample space Ω , specification that a trial of kind S & T is occurring at t on the chance setup in question, and, finally, the assumption that the information that the trial is of kind T is stochastically irrelevant information.

This is precisely the kind of information presupposed for the purpose of direct inference from composite chance hypotheses in sections 12.13–12.16. Its centrality in discussions of inverse inference from knowledge of test behavior to credal states over rival statistical hypotheses emphasizes the fact that knowledge of chances cannot be obtained from the testimony of the senses alone. As noted before, expansion on the basis of the data to a conclusion rejecting some of the rival simple statistical hypotheses depends on a successful use of the data to obtain suitable posterior credal distributions over the rival alternatives; and, as I have been stating and restating, such use of Bayes' theorem depends on likelihood functions which, if they are to be nonproblematic consequences of the use of direct inference, presuppose substantial background information.

Thus, even that form of inverse inference to statistical hypotheses involving the assignment of credal probabilities to hypotheses on the basis of data presupposes that X 's corpus prior to experimentation extends substantially beyond truths of logic and set theory, alleged conceptual truths, the testimony of the senses, and the records of the memory. If the likelihood functions on the data are uniquely determined by the principle of direct inference and the background knowl-

edge, that background knowledge will contain substantial information about chances.

Such knowledge is clearly open to revision. Some information about chances may be removed and other information added in its stead. But if such revisions are subject to critical control and a systematic account of the matter is worth constructing, the problem of inductive expansion or inductive acceptance must be recognized as an important problem for epistemology and especially that part of epistemology concerned with the conditions under which experimental data can be used to determine credal probability assignments to statistical hypotheses via Bayes' theorem.

13.4 Unbiased Priors

Let K be X 's initial corpus relative to which at least and at most one hypothesis h_θ where $\theta \in \Theta$ (or $h_\theta \in U$) is true and each is consistent with K . Each of the h_θ 's is a simple chance hypothesis specifying a chance distribution over the sample space Ω on trials of kind S on chance setup a . K also contains information to the effect that at some specific time a trial of kind S is to occur which is also a T and that the information that the trial of kind S is also of kind T is stochastically irrelevant. If e_α asserts that on the trial in question, result of kind R_α occurs, the principle of direct inference mandates that every permissible Q -function in $C(K)$ satisfy the requirement that $Q(e_\alpha; h_\theta) = F_\theta(\alpha)$. Thus, direct inference guarantees the existence of a unique likelihood function on the data e_α . The anxieties of the previous section are allayed, at least for awhile.

X presumably begins with some sort of confirmational commitment just as he begins with an initial corpus. S. Spielman has raised doubts about this.⁷ His doubts would be entirely justified were X obliged to have a numerically precise confirmational commitment specifying a uniquely permissible Q -function for every potential corpus.

The view I favor does not demand this. Indeed, X at the outset could embrace a confirmational commitment that allows every Q -distribution over the elements of U to be permissible conforming to the requirements of the calculus of probability—provided, of course, that these distributions also conform to the requirements of inductive logic, confirmational condition-

alization, credal convexity, and the constraints otherwise adopted by X for his confirmational commitment.

I suspect that in many experimental contexts X might very well begin with such a maximally indeterminate prior—although this need not always be so. For the present, however, I shall suppose that he does.

But as we have seen in section 13.2, such maximal indeterminacy implies that the data obtained as a result of experimentation will fail to lead to any change of credal state if the confirmational commitment remains unchanged. Objectivist necessitarians obligate X to retain the confirmational commitment unrevised. However, once necessitarianism is rejected, revisions of confirmational commitment are allowed provided that there is good reason for such revision. It is clear that if the prior credal state is initially maximally indeterminate, there is very good reason for attempting to strengthen the confirmational commitment by removing some of the prior Q -distributions over the elements of U .

The problem is, however, how such strengthening should be done. There should be some good reason for strengthening one way rather than another. If there is none, one should suspend judgment between all allowable strengthenings through adopting their convex hull as the prior credal state. If no potential strengthening is ruled out, this implies that no strengthening will take place.

We cannot, however, appeal to considerations of inductive logic to help us in establishing the superiority of one strategy for strengthening over another.

However, if X is conducting the experiment in order to find out which element of U is true and if he has evaluated the rival potential answers with respect to informational value and has adopted a value q for his index of boldness (or caution), X might reasonably argue that any Q -distribution over the elements of U which is such that, for some h_θ , $Q(h_\theta) < qM(h_\theta)$ ought to be ruled out of consideration as the uniquely permissible element of the revised prior credal state. To strengthen to such a prior credal state would justify X in rejecting elements of U , whereas prior to the strengthening he was not entitled to do that.

Any strengthening leading to such a result is objectionable; for X would then rationalize an expansion without changing

the evidence or the demands for information, but merely by altering his prior credal state.

In cases where U is finite, any prior Q -distribution for which $Q(h_\theta) \geq qM(h_\theta)$ avoids the objectionable consequence and, hence, is a candidate for consideration for the purpose of strengthening the prior credal state. But if there is no basis for discriminating between such *unbiased* Q -distributions, none of them should be ruled out and the revision of the prior credal state should involve shifting to the set of unbiased prior Q -distributions relative to the corpus K , ultimate partition U , information-determining M -function (section 2.4), and index q . Notice that if $q = 1$, there will be exactly one unbiased prior Q -distribution—namely, $Q(h_\theta) = M(h_\theta)$.

If there is more than one permissible M -function according to X 's demands for information, the prior credal state should consider the set of all unbiased priors relative to all of these M -functions (which form a convex set).

13.5 Strongly Un- biased Priors

When U is infinite, special technical problems arise.

Let U be countably infinite and suppose each element of U is informationally as valuable as any other. The M -function must then assign $M(h_\theta) = 0$ for all $h_\theta \in U$.

According to my rejection rules, an element of U is to be rejected relative to whatever corpus is given if and only if $Q(h_\theta) < qM(h_\theta)$. However, in this case, no such rejection is ever possible, since $M(h_\theta) = 0$. Furthermore, we cannot seek relief from the difficulty by shifting to densities as we can when we have continuous parameters and the M -distribution is characterizable by a density function. Obviously the methods for evaluating expansion strategies have to be revised.

My proposal is this. Take any element h_θ in U . Take any set of n elements of U including h_θ and find the conditional Q -distribution over elements of this subset given that one of them is true. The corresponding conditional M -distribution assigns each element in the subset an $M > 0$. Indeed, when the M -function is uniform, then $M = 1/n$ for each element in the subset.

Apply the usual rejection rule and reiterate the application until a stable result is obtained in the sense of section 2.8. If there is no set of n elements of U containing h_θ leading to

stable rejection of h_θ , then h_θ is *n-unrejected*. If there is such a set, it is *n-rejected*.

h_θ is rejected if and only if it is *n-rejected* for every n greater than some n^* .

A Q -distribution over U (or the σ -finite measure representing it) is unbiased relative to K , U , M , and q if and only if for any h_θ in U there is no n^* such that h_θ is *n-rejected* for every $n > n^*$.

Notice, however, that unbiasedness does not preclude an element of U from being rejected relative to the Q -distribution for some n -member subset of U . It seems reasonable, however, in this case that we would seek to prevent any element of U from being *n-rejected* for any positive value of n prior to experimentation. A distribution over countably many elements of U having this property is *strongly unbiased*.

In the countably infinite case, we have thus far defined strong unbiasedness only when the M -function is uniform. However, the characterization applies just as well to situations where the M -distribution is not uniform.

Furthermore, strong unbiasedness can be defined for cases where U is finite as well.

In both the countably infinite and the finite case, a strongly unbiased prior is such that for no finite subset of U in which h_θ belongs is h_θ rejected given that some element of that subset is assumed true but otherwise matters are left as they were.

Finally, if θ is continuous and the information-determining M -function is represented by a density, we can require for strong unbiasedness that if we consider a set of alternatives of the form "the true value of θ falls in the interval from θ_i to $\theta_i + d\theta$ " for finitely many distinct θ_i 's and conditionalize the Q -function and M -function on the assumption that one of these alternatives is true, then no matter how small but positive $d\theta$ is allowed to be, no element of the set of alternatives is rejected for the given value of q .

I propose that when X begins with a situation where the prior credal state for U is maximally indeterminate and X seeks to strengthen it, he should adopt as his new prior credal state the set of all strongly unbiased Q -functions relative to K , U , q , and every permissible M -function according to his demands for information.

13.6
Standardized
Priors and
Posteriors

Let us focus on cases where there is exactly one M -function in the set representing X 's demands for information.

When $q = 1$, there can be at most one strongly unbiased prior Q -distribution $Q^*(h_\theta) = M(h_\theta)$ relative to U and M .

Even when $q < 1$, Q^* is strongly unbiased. I shall call Q^* the *standardized prior Q -distribution*.

Corresponding to the standardized prior distribution, there is a standardized posterior distribution $Q^*(h_\theta; e_\alpha)$ conditional on e_α .

Letting the likelihood $L(h_\theta; e_\alpha) = Q(e_\alpha; h_\theta)$ be given, let

$$L^*(e_\alpha) = \sum_{\theta \in \Theta} L(h_\theta; e_\alpha)M(h_\theta).$$

Then

$$(1) \quad Q^*(h_\theta; e_\alpha) = L(h_\theta; e_\alpha)M(h_\theta)/L^*(e_\alpha).$$

Whatever the index of caution q might be, if rejections are based on the standardized posterior, h_θ is rejected if and only if

$$(2) \quad Q^*(h_\theta; e_\alpha) < qM(h_\theta).$$

In the light of (1), (2) holds if and only if

$$(3) \quad L(h_\theta; e_\alpha) < qL^*(e_\alpha).$$

When $q = 1$, condition (3) states a necessary and sufficient condition for rejecting an element of U on the data e_α assuming that X began with a prior credal state consisting of strongly unbiased priors.

In that special case, reiterating the rejection rule until a stable conclusion is obtained yields the result that all elements of U are rejected but those bearing maximum likelihood on the data.

In most contexts, it seems to me that X should not be so bold as this but should adopt a $q < 1$.

In that case, the prior and the posterior credal states will be indeterminate. The rejection rule stipulates that h_θ be rejected if and only if the maximum value for $Q(h_\theta; e_\alpha)$ derived by Bayes' theorem from a strongly unbiased prior is less than $qM(h_\theta)$.

To compute exact values for these maximum posterior values will be difficult in many cases. However, it is possible to

obtain approximations for such maxima which are multiples of the standardized posterior values. In some cases, the approximation is rather good. In other cases, it is not so good. However, the method of approximation proposed has the virtue that when applied to the problem of rejection it errs on the side of caution. It leads to rejecting only elements of U that would be rejected utilizing exact methods, although it sometimes fails to lead to rejections permitted by exact methods.

13.7
Stable
Estimation

Take the easiest case first. Suppose U is finite and the M -function assigns the same M -value of $1/n$ to each h_θ in U .

Take any value θ of the parameter and consider the strongly unbiased prior distribution over U assigning h_θ a Q -value which is a maximum assignable by any strongly unbiased prior.

If $h_{\theta'}$ is any other element of U , it is known that

$$Q(h_\theta; h_\theta \vee h_{\theta'}) \leq 1 - q \frac{M(h_{\theta'})}{M(h_\theta \vee h_{\theta'})} = 1 - \frac{q}{2} = \frac{2 - q}{2}.$$

According to that Q -function, $Q(h_{\theta'}; h_\theta \vee h_{\theta'}) \geq q/2$.

From this it follows that

$$Q(h_\theta) \leq Q(h_{\theta'}) \frac{(2 - q)}{q}.$$

This is true for every $\theta' \neq \theta$. Hence, to obtain a maximum prior Q -value for h_θ , every other element of U should be assigned the same Q -value x and $Q(h_\theta)$ should equal $(2 - q)x/q$.

Since the number of elements of U is a finite number n , $x = q/[q(n - 2) + 2]$.

Since the M -function assigns equal M -value to all elements of U ,

$$L^*(e_\alpha) = \sum_{\theta' \in \Theta} L(h_{\theta'}; e_\alpha)M(h_{\theta'}) = \sum_{\theta' \in \Theta} \frac{L(h_{\theta'}; e_\alpha)}{n}$$

(see section 13.6).

Let

$$L_\theta^*(e_\alpha) = nL^*(e_\alpha) + L(h_\theta; e_\alpha) \frac{(2 - 2q)}{q}.$$

If every $h_{\theta'}$ in U distinct from h_{θ} is assigned the prior Q -value x and h_{θ} is assigned the prior Q -value $(2 - q)x/2$, the posterior value for h_{θ} determined by Bayes' theorem must be

$$(1) \frac{L(h_{\theta}; e_{\omega})(2 - q)}{L_{\theta}^*(e_{\omega})q}$$

(1) must be the maximum posterior value for h_{θ} allowed by the given likelihood function, Bayes' theorem, and the set of strongly unbiased priors. But (1) is less than

$$(2) \frac{L(h_{\theta}; e_{\omega})(2 - q)}{nL^*(e_{\omega})q} \leq \frac{L(h_{\theta}; e_{\omega})(2 - q)}{L^*(e_{\omega})q} = Q^*(h_{\theta}; e_{\omega}) \frac{(2 - q)}{q}$$

$Q^*(h_{\theta}; e_{\omega})$ is the standardized posterior of section 13.6.

Thus, we may use (2) as an approximation for the maximum posterior value for h_{θ} obtainable from a strongly unbiased prior. This approximation becomes exact when $q = 1$ and is excellent for values of q near 1. Moreover, even when q is fairly small, when $L^*(e_{\omega})$ is large compared to $L(h_{\theta}; e_{\omega})$ the approximation will continue to be a good one.

These results extend quite naturally to cases where U is countably infinite and the uniform M -function is represented by a σ -finite measure. It also can be extended to cases where θ is continuous and the M -function representable by a continuous density. In this guise, the claims just made are known as theorems of "stable estimation."⁸

Theorems of stable estimation apply to cases where priors are approximately uniform—i.e., where they are strongly unbiased relative to uniform M -functions. It is of interest to extend these theorems to cover cases where posteriors are derived from approximations of standardized priors—i.e., priors strongly unbiased relative to other M -functions besides uniform ones.

Consider the finite case. Order all elements of U in order of increasing M -value. The order of elements of U with equal M -value is immaterial.

Let

$$Z(\theta_i, \theta_{i'}, q) = \begin{cases} \frac{M(h_{\theta_i}) + (1 - q)M(h_{\theta_{i'}})}{q} & \text{if } i \neq i' \\ M(h_{\theta_i}) & \text{if } i = i'. \end{cases}$$

Let Q be the Q -function assigning maximum value to h_{θ_i} among all strongly unbiased priors. Then

$$Q(h_{\theta_i}) \geq \frac{Q(h_{\theta_i})M(h_{\theta_i})}{Z(\theta_i, \theta_j, q)}$$

for every h_{θ_j} where $j \neq i$. If $i = j$, the equality is strict.

Let

$$\begin{aligned} L_{\theta_i}^{**}(e_{\omega}) &= \sum_{j=1}^n L(h_{\theta_j}; e_{\omega})Q(h_{\theta_j}) \\ &\geq \frac{\sum_{j=1}^n L(h_{\theta_j}; e_{\omega})Q(h_{\theta_j})M(h_{\theta_j})}{Z(\theta_i, \theta_j, q)} \end{aligned}$$

Let j^* be such that $Z(\theta_i, \theta_{j^*}, q)$ is the maximum for all Z -values where θ_i is held fixed. Then

$$L_{\theta_i}^{**}(e_{\omega}) \geq \frac{\sum_{j=1}^n L(h_{\theta_j}; e_{\omega})Q(h_{\theta_j})M(h_{\theta_j})}{Z(\theta_i, \theta_{j^*}, q)} = \frac{L^*(e_{\omega})Q(h_{\theta_i})}{Z(\theta_i, \theta_{j^*}, q)},$$

and

$$\begin{aligned} Q(h_{\theta_i}; e_{\omega}) &= \frac{L(h_{\theta_i}; e_{\omega})Q(h_{\theta_i})}{L_{\theta_i}^{**}(e_{\omega})} \\ &\leq \frac{L(h_{\theta_i}; e_{\omega})Z(\theta_i, \theta_{j^*}, q)}{L^*(e_{\omega})} \\ &= \frac{L(h_{\theta_i}; e_{\omega})M(h_{\theta_i})Z(\theta_i, \theta_{j^*}, q)}{L^*(e_{\omega})M(h_{\theta_i})} \\ &= \frac{Q^*(h_{\theta_i}; e_{\omega})Z(\theta_i, \theta_{j^*}, q)}{M(h_{\theta_i})} \end{aligned}$$

We know already that when $q = 1$, $Q(h_{\theta_i}; e_{\omega})$ should equal the standardized posterior $Q^*(h_{\theta_i}; e_{\omega})$. Compatibility with this requirement is guaranteed by the fact that the quantity $Z(\theta_i, \theta_{j^*}, q)/M(h_{\theta_i})$ approaches 1 as q approaches 1.

As $M(h_{\theta_i})$ approaches $M(h_{\theta_{j^*}})$ the ratio approaches $(2 - q)/q$.

These results may be extended to cover cases where n goes to infinity, and continuous cases as well.

13.8

Likelihood and Rejection Rules

This discussion began by considering a situation where X is concerned to conduct an experiment and obtain data on the basis of which elements of some subset of the ultimate partition U will be rejected. X has demands for information rep-

resented by an information-determining M -function and is committed to a standard of caution q .

Before conducting the experiment and relative to the initial corpus K , X 's prior credal state for elements of U is maximally indeterminate. This circumstance furnishes a warrant for X 's revising his confirmational commitment by strengthening his prior credal state for elements of U in order that the data of experiment will be worth acquiring for the purposes of the inquiry being undertaken.

In this kind of situation, X should shift to the prior credal state consisting of all strongly unbiased prior distributions over elements of U relative to K , U , M , and q .

In section 13.7, methods for approximating maximum permissible Q -values for elements of U conditional on e_α were introduced; it was there shown that for a given h_θ the approximation should be some multiple of the standardized posterior $Q^*(h_\theta; e_\alpha)$.

It is now time to explore the application of these approximate maximum posterior values to the problem of inferential expansion when X finds out that e_α is true after experimentation.

In the special case where the demands for information are represented by a uniform M -distribution, it has been shown that every permissible posterior Q -distribution is such that

$$Q(h_\theta; e_\alpha) < \frac{Q^*(h_\theta; e_\alpha)(2 - q)}{q}.$$

Furthermore, h_θ in U is to be rejected if and only if

$$Q(h_\theta; e_\alpha) < qM$$

for every permissible posterior Q -function, where $M = M(h_\theta) = 1/n$ for all h_θ in U .

Consequently, if we are content with approximation which errs on the side of caution (i.e., leads to rejecting fewer elements of U than exact methods warrant), we should adopt the prescription recommending the rejection of h_θ in U if and only if

$$Q^*(h_\theta; e_\alpha) < \frac{q^2 M}{2 - q}.$$

But this condition holds if and only if

$$L(h_\theta; e_\alpha) < \frac{q^2 L^*(e_\alpha)}{2 - q},$$

where

$$L^*(e_\alpha) = M \sum_{\theta \in \Theta} L(h_\theta; e_\alpha),$$

as section 13.6 requires.

This approximate rule recommends rejecting an element of U if and only if its likelihood on the data is less than $L^*(e_\alpha)$ times some constant $k \leq 1$. The value of the constant k varies from 1 when $q = 1$ to 0 when $q = 0$.

If the rejection rule is reiterated until a stable result is obtained, the stable result specifies that an element of the initial U is rejected if and only if its likelihood on the data is less than some fraction of the maximum likelihood on the data.

Both the initial likelihood-based rejection rule and the one derived from reiteration have been discussed in the literature before.⁹ What is interesting about the argument for utilizing such rules here is that, according to the approach adopted, likelihood-based rejection rules are not taken to be fundamental but to be applicable when the circumstances are appropriate.

Furthermore, the relevant contextual parameters considered here concern not only the background information K and the initial maximally indeterminate prior credal state, but also the ultimate partition U , the demands for information that rated all elements of U equally informative, and the value of the index q .

The conditions for reaching the results described will not always be satisfied. Likelihood-based rejection rules cannot be used in inductive expansion on every occasion. The context matters. What has been shown is that certain combinations of contextual factors do justify the use of likelihood-based rejection rules.

Furthermore, it is entirely plausible to suppose that in a great many contexts of research where the potential answers are rival simple statistical hypotheses or alternations of such, the information-determining M -function will rate all simple alternatives as equally informative. There is no principle of

reason which mandates this. The demands for information occasioning the inquiry will control the M -function. Perhaps these demands themselves are determined by research programs of some sort. We need not settle that matter here. I contend only that on many occasions rival simple statistical hypotheses will be counted as equally informative.

It is in inquiries of this kind that prior credal states are often plausibly strengthened to sets of strongly unbiased priors which approximate uniform priors. As we have seen, the posterior distributions conform approximately to the requirements of the likelihood function. This circumstance explains the usefulness of likelihood-based rejection rules.

Because the circumstances for legitimately using likelihood-based rejection rules may be fairly widespread, there is pre-systematic cogency to adopting likelihood-based rules that renders it tempting to try to make likelihood a fundamental notion in inference and to take likelihood-based rejection rules as universally applicable.

In my opinion, this is a mistake. Prior credal states should not always be unbiased; and, even when there is justification for endorsing such states, the M -function need not always be uniform. When it is not, approximate rejection rules are no longer always likelihood-based in the sense explained previously.

Suppose the M -function defined over finite U is not uniform. The rejection rule stipulates that h_θ is rejected, given the new data e_α , if and only if $Q(h_\theta; e_\alpha) < qM(h_\theta)$ for every permissible Q -function.

According to section 13.7,

$$Q(h_\theta; e_\alpha) \leq \frac{Q^*(h_\theta; e_\alpha)Z(\theta, \theta^*, q)}{M(h_\theta)}.$$

The right-hand side of this inequality is less than $qM(h_\theta)$ if and only if

$$L(h_\theta; e_\alpha) < \frac{qM(h_\theta)L^*(e_\alpha)}{Z(\theta, \theta^*, q)}.$$

This latter condition may serve as an approximate rejection rule when the M -function is not uniform. The right-hand side depends on the value of θ when the M -function is not uniform and q is less than 1. Hence, when the M -function is not uni-

form and $q \ll 1$, the rejection rule does not stipulate, even as an approximation, that h_θ is to be rejected if and only if its likelihood on the data is less than a fixed level which is the same for all elements of U . The rejection level depends on the value of $M(h_\theta)$.

13.9 Ignorance and Strongly Un- biased Priors

Since Bayes and Laplace, when X is in some sort of state of ignorance over elements of U , he has been urged to adopt as his prior credal state a uniform distribution over the alternatives. This recommendation is often advocated as being grounded on context-independent principles of inductive logic which mandate what every rational X should adopt as his credal state given that his corpus of knowledge represents a suitable sort of ignorance.

Principles of insufficient reason are plagued with all sorts of difficulties. Inconsistencies threaten and ad hoc repairs are made to prevent trouble, thereby removing whatever shred of presystematic cogency the principles might have had in the first place. Both proponents of such principles and many of their critics are concerned with the objectivity of the prior distributions used in deriving posteriors. The critics rightly complain of the arbitrariness involved in the choice of sets of alternative hypotheses over which principles of insufficient reason are to be applied.

Critics tend to respond in one of two ways. They either become personalists and adopt priors without any justification other than the alleged need to adopt some numerically precise one. Or they refuse to adopt any precise prior at all, becoming, on my reconstruction, advocates of maximally indeterminate priors.

I have been advocating a point of view according to which considerations other than logic, language, and the available evidence play a role in the justification for adopting a prior credal state. The problem under investigation, the rival potential answers identified for it, their evaluation with respect to informational value, and the index of caution may also play a role.

Often when such factors are taken into account, the prior credal state that results consists of a family of priors resembling those the advocates of insufficient reason recommend.

But there are some vital differences. First, only rarely will

the prior recommended be precise. Second, the prior is not recommended by an appeal to principles of inductive logic applicable regardless of the question under consideration and the demands for information involved in it.

Critics may complain that the view which emerges is no different from the personalist position. Insisting on the context-dependence of choices of prior credal states is no protection against subjectivism.

Suppose X begins with the same initial corpus and ultimate partition but alters his M -function or his index q . The set of strongly unbiased priors will vary in ways which could lead to markedly different posteriors and rejections.

Of course, X might change his M -function arbitrarily in the manner indicated. But he might change his corpus arbitrarily too. The fact that credal states are dependent on X 's body of knowledge is often acknowledged, even though this is a contextual and historical factor just as X 's demands for information are.

Needless to say, X should not alter either his corpus or his demands for information willfully. Changes in corpus and in demands for information ought to be kept under careful critical control or, at any rate, within certain limits they should be.

It is true that the revisionist approach I favor agrees with both intemperate and tempered personalism in giving up the pipe dream of context-independent objectivity in inquiry. But revisionism does not sanction anarchy. Nor does it maintain that probability judgment depends on the context and leave it entirely to judgment and insight to say how context controls probability judgment.

Return to the problem discussed by Fisher of testing a black mouse to find out whether it is homozygous or heterozygous. The experiment is to mate it with a brown mouse and note the color of the offspring.

We shall consider the case where no information about the origin of the mouse is available so that, in Fisher's words, "no experimenter would feel he had warrant for arguing as if he knew that of which in fact he was ignorant."

I take this to mean that no experimenter begins his investigation with a numerically definite prior obtained via direct inference from knowledge of chances. No such knowledge is available. Hence, at the outset, the investigator X adopts a

maximally indeterminate prior credal state for the two rival hypotheses (h_1 asserting homozygosity and h_2 asserting heterozygosity).

In this case, I take it that most investigators would agree that the two rival hypotheses bear equal M -value of $.5$. If they do not, they would have to consider all seriously proposed M -values and take their convex hull. But I shall simplify by supposing that there is exactly one permissible M -function which assigns both alternatives equal M -value.

In this case, the class of strongly unbiased priors permits values for $Q(h_1)$ to range from $.5q$ to $1 - .5q$ and similarly for $Q(h_2) = 1 - Q(h_1)$.

If the test mouse is mated with a brown mouse (known, therefore, to be homozygous) and h_1 is true, the chance of obtaining only black offspring is 1. If h_2 is true, the chance of r blacks out of n offspring is $\binom{n}{r}(.5)^n$.

Clearly, if r is less than n , the hypothesis h_1 is false and h_2 is established. But suppose all n offspring are black. The chance of this happening if h_2 is true is $(.5)^n$. Let this result be reported by e_n . $Q(e_n; h_1) = 1$ whereas $Q(e_n; h_2) = (.5)^n$.

The maximum posterior for h_2 is $\frac{(.5)^n(1 - .5q)}{(.5)^n(1 - .5q) + .5q}$.

For h_2 to be rejected, this ratio would have to be less than $.5q$.

The maximum posterior for h_1 is $\frac{1 - .5q}{1 - .5q + (.5)^{n+1}q}$.

The rejection condition is, once more, that this be less than $.5q$.

It is clear that if $q = 1$, then h_2 is rejected because its likelihood on the data is below the threshold. On the other hand, as q is reduced, h_2 may escape rejection unless n is increased. The hypothesis h_1 , of course, avoids rejection because its likelihood always remains higher than that of h_2 .

When $q = 1$, a numerically precise and uniform prior should be adopted as proponents of insufficient reason advocate. However, the experimenter should not adopt this prior by arguing "as if he knew that of which he is in fact ignorant." He should adopt the uniform prior because it is the only prior distribution which avoids prejudging the conclusion to be adopted via inferential expansion.

If $q < 1$, there are many such distributions and X is not

entitled to select one of them in preference to the others. He should suspend judgment between them all.

If $q = 1$, observation of one black offspring (and no brown) suffices to warrant rejection of h_2 . If $q = .5$, it takes observation of four black offspring to warrant such rejection. If $q = .05$, observation of seven is needed.

These calculations are exact. The approximate methods discussed before favor rejecting h_2 if and only if

$$L(h_2; e_n) = .5^n < (1 + .5^n) \cdot 5 \frac{q^2}{2 - q}.$$

As before, if $q = 1$, it takes one black offspring to warrant rejection of h_2 and four if $q = .5$. However, if $q = .05$, 11 black offspring are needed.

The approximate methods do not err at all until q becomes less than $.5$; and when they do, they err on the side of caution.

A. Shimony appealed to the demands of the problem under investigation and the potential answers identified as solutions to that problem in considering how prior credal probability judgments are to be modified.¹⁰ He too was concerned that priors be modified so that seriously proposed potential answers be capable of "winning" or "losing" in some sense as a result of inquiry but without prejudicing the result.

My proposals are made in the same spirit. They too are tied to the problem under investigation and potential answers identified by the inquirer.

Unlike Shimony, I construe "winning" and "losing" in terms of inferential expansion—i.e., whether a hypothesis is admitted into evidence, its negation is, or neither. Furthermore, although Shimony acknowledges that his methods will not rule out all but one prior distribution, he still urges picking one of these as uniquely permissible. He remains a personalist—albeit a tempered personalist.

These differences ought not, however, to obscure the important points of agreement between Shimony's view and mine.

Contextual considerations may be systematically exploited without breeding rampant psychologism or constructing fantastic third worlds. We need not be burdened by the curse of Frege—who has frightened so many with the threat of psychologism that they retreat to a sterile pseudo objectivity,

leaving many important matters open to historical, psychological, and sociological study but immune to critical review.

13.10 Prior Data

Many occasions arise where the conditions for applying the methods for selecting prior credal states just outlined do not apply.

For example, when prior credal states are derivable via direct inference from knowledge of chances already available, principles of inductive logic mandate the choice of a fairly determinate credal state prior to experimentation.

But even when inductive logic fails to determine a prior credal state relative to prior background knowledge, X might already have some confirmational commitment which assigns the elements of U a fairly determinate prior credal state.

Suppose X is convinced that the black test mouse is the offspring of a homozygous black and a homozygous brown mouse. Given his knowledge of genetics, X knows the test mouse is heterozygous. Hence, he is not at all concerned to test the mouse for heterozygosity. Why should he do so when he already knows the answer? Perhaps, however, he has some other reason for mating the mouse with a brown mouse (known to be homozygous). He might be interested, for example, in illustrating some aspect of the workings of Mendelian genetics to students.

In particular, he might be interested in showing the students that the mouse's offspring will be approximately 50% black and 50% brown as Mendelian theory predicts.

Suppose that 20 offspring are obtained and they are all black. The result is perfectly consistent with the assumption in X 's initial corpus K_1 that the mouse is heterozygous. Yet, on the assumption, the result is not "what is to be expected." X might wish, therefore, to give a hearing to a rival hypothesis which can do better in explaining the result.

He could question Mendelian theory; but the loss of information incurred would be severe. It would be preferable to question the parentage of the mouse and contract from the corpus K_1 (the expansion of K_1 by adding the data concerning the 20 offspring) to K_2 , relative to which X suspends judgment between the hypothesis of homozygosity h_1 and of heterozygosity h_2 .

Notice that the data concerning the 20 offspring are not

evidence against heterozygosity when K_1 is the corpus. To say so is incoherent; for the assumption of heterozygosity is part of the evidence.

The data provide a good reason for contracting the corpus K_1 to K'_2 so that the hypothesis h_1 of homozygosity can be given a hearing (see chapter 3).

Let K'_2 be the new contracted corpus. X 's credal state for rival hypotheses about genotype of the test mouse could be maximally indeterminate. Were this the case, X 's confirmational commitment would have to be such that the credal state $C(K_2)$ for h_1 and h_2 would also have to be maximally indeterminate where K'_2 is the expansion of K_2 obtained by adding the data about the 20 mice to K_2 .

Clearly, however, X might have already been committed to a confirmational commitment according to which the distributions for h_1 and h_2 in $C(K_2)$ are strongly unbiased relative to K_2 , U , M , and q (where M assigns the two elements of U equal value of .5). The set of distributions over the elements of U in $C(K_2)$ will no longer be indeterminate. Indeed, the maximum Q -value for h_2 will be sufficiently low to reject the hypothesis of heterozygosity without any further experimentation. The observation of the 20 black mice would be evidence enough to warrant the conclusion.

There are other scenarios to consider; but we cannot exhaust them all here. Let us focus on the two just mentioned.

In the case where the credal state for elements of U relative to K'_2 is indeterminate, the analysis offered in the previous sections is applicable. X should strengthen to the set of strongly unbiased distributions *relative to* K'_2 . He should not strengthen to the set of strongly unbiased distributions relative to K_2 ; for he would then, by a change in confirmational commitment alone, so change his appraisals of risk as to warrant expansion which was not justified initially. Such shifts reveal, in my opinion, a lack of respect for the desirability to avoid error or, at least, to take risk of error seriously.

Thus, in the case under consideration, X will be obliged to conduct additional experiments to determine the fate of the two hypotheses about the genotype of the test mouse. The data concerning the 20 mice were warrant for opening up a new inquiry but were insufficient warrant for then replacing

the hypothesis of heterozygosity with the hypothesis of homozygosity.

Turn now to the second case, where X has an antecedent confirmational commitment which assigns h_2 extremely low credal probability relative to K'_2 .

There are two possibilities. X might be prepared to expand by rejecting H_2 without further experimentation. In that event, the data about the 20 mice not only justified contraction from K_1 to K'_2 , but also the subsequent expansion to the conclusion that the mouse is homozygous.

But X might demur. He might consider that he had contracted initially in order to give the hypothesis of homozygosity a hearing. But there are two senses to giving that hypothesis a hearing. One can give a hearing relative to the contracted corpus or one can give a serious hearing through an inquiry which begins with the contracted corpus but then obtains new information (through experimentation and other side investigations) that is then to be used to decide the issue.

I know of no general principle which can be used to decide when one should rest content with giving a mere hearing or whether one should give a serious hearing. I suspect that this matter is resolved by considering a trade-off of several epistemic benefits.

In any case, to give a fair or serious hearing involves preventing a resolution of the matter prior to further experimentation.

(Incidentally, there is an interesting difference between cases where upon contraction there is a warrant for expansion reinstating the hypothesis initially removed and cases where there is a warrant for expansion introducing a rival to that hypothesis as in the example considered here.)

In our example, giving such a hearing requires modifying the prior credal state. However, such modification ought not ignore the initial credal state. X , after all, begins with an earnest commitment to it. Claims are made that he should modify his commitment so as to prevent expansion by accepting h_1 without further experimentation; but this can be done without leading X to rule out the distributions he initially considers permissible as being impermissible. To the contrary, he can retain these distributions as permissible and add others to prevent a premature resolution of the issue. He can, in

particular, adopt the convex hull of the set of distributions he initially regards as permissible and the set of strongly unbiased distributions.

In our example, doing so will lead to a situation little different in practice from one where the prior credal state consists exclusively of strongly unbiased priors. This is because the only circumstance under which h_1 can be rejected is where a sample containing at least one brown mouse is obtained.

However, cases arise where the rival hypotheses are all such that rejection does not follow from deductive logic and the data alone. The methods just illustrated, when applied to such cases, will give an edge to hypotheses strongly probable according to the initial confirmational commitment by rendering them more difficult to eliminate on the basis of the data. Yet, it will not be impossible to eliminate them. The other alternatives will receive a hearing, resulting in a demand for the acquisition of new data.

These considerations illustrate further how contextual considerations can lead to the revision of confirmational commitments. They also indicate that the task of exploring the ways in which confirmational commitments may legitimately be revised is important unfinished business.

13.11 Irrelevance Revisited

Part of the unfinished business, it should not be forgotten, is the determination of likelihood functions. Such determination depends, of course, on the principle of direct inference. But it depends also on knowledge of chances sufficient to secure stochastic irrelevance of extra information about kinds of trials. If such knowledge is lacking, X might sometimes endorse a confirmational commitment according to which the extra information about the kind of trial is confirmationally irrelevant to hypotheses about the outcome of testing.

An important question about the revision of confirmational commitments concerns conditions under which X is justified in modifying his confirmational commitment so as to secure such a result.

The problem is a difficult one which deserves further investigation. In practice, it can sometimes be avoided by redesigning experiments so that one can rely on background knowledge of stochastic irrelevance. But this is not always feasible; and randomization is not the panacea Fisher alleged it to be.

13.12 Probability and the Growth of Knowledge

In chapters 1-3 of this book I have argued for the importance of developing an account of the improvement of knowledge construed as a standard for serious possibility used as a resource in inquiry and deliberation.

In chapters 4-13 I have sought to extend the outlook of the first three chapters to the revision of probability judgment. X 's cognitive resources for inquiry and deliberation include both his standards for serious possibility and his appraisals of hypotheses counted as seriously possible with respect to credal probability.

I have not pretended here to offer a complete account of either the improvement of knowledge, confirmational commitments, or credal states. But I have offered a framework for discussing these questions and made a few substantive claims.

Several large conclusions may be drawn if the approach advocated here is on the right track. (1) Contextual considerations can be taken into account in a systematic manner without precluding the critical control of the revision of knowledge. (2) Those authors who deny the existence of a fixed scientific method are, in all essentials, in the right. The principles of inductive logic, deductive logic, rational choice, and valuation proposed here are relatively weak and are applicable to all aspects of conduct. Any stronger canon of scientific method that may prevail at some historical moment is subject to critical review. The growth of knowledge is at the same time the development of method. (3) The problem of acceptance will not go away no matter how often simple-minded Bayesians may declare it dead. The question of revising probability judgments and the question of revising standards for serious possibility (which is the problem of acceptance) are inextricably linked with one another. (4) There is no gulf between the weak principles of rationality regulating practical deliberation and theoretical inquiry. (5) On the other hand, the cognitive values of scientific inquiry have a life of their own distinct from the values of commerce or politics, or from ethical considerations.

Even those who disagree with some or all of these large conclusions or with some of the technical proposals made here may, nonetheless, find some value in some of my proposals. The account of indeterminate probability that has been constructed is useful as a means for comparing rival views of

probability judgment and induction without begging questions for or against competing views.

To demonstrate this, the remaining chapters of this book will be devoted to discussing some views concerning probability judgment alternative to the one I favor. These views are all, in a sense, understandable as responses to the difficulties facing objectivist necessitarianism that attempt to meet the problems while somehow remaining loyal to the spirit if not the letter of objectivist necessitarianism.

The discussion will, I hope, illustrate the use of the methods developed here for representing probability judgments. It will also contribute to an important step of my argument which has not, as yet, been taken.

In this chapter I have assumed that objectivist inductive logic is a complete inductive logic but that objectivist necessitarianism is not viable. On this basis, I have concluded that some sort of revisionism ought to be favored.

I do not know how to prove that objectivist inductive logic is complete. However, we can consider ways to improve on objectivist necessitarianism to see whether there is a promising alternative to revisionism. I do not think that there is; but there is a need to consider the chief alternatives.

14.1 Introduction

In 1930, R. A. Fisher proposed a form of argument which he seems to have thought would yield numerically definite posterior distributions without commitment to numerically definite priors.¹ Fisher called such arguments "fiducial" inferences. In his later writings, Fisher explicitly insisted that fiducial probability distributions are appraisals of hypotheses with respect to probability on the data—as are appraisals on the data derived via Bayes' theorem—even though in fiducial inference there is no appeal to prior distributions.²

If Fisher is to be taken seriously, he must have been committed to rejecting one of the conditions which have been imposed on confirmational commitments in this book; for these conditions imply that no numerically definite posterior can be obtained unless there is a numerically definite prior.

Fisher failed to indicate which of the principles proposed previously he was prepared to modify. He appears to have invoked all of them at least tacitly at one point or another. Hence, it is almost impossible to make an educated guess as to what his considered opinion was.

Nonetheless, I am inclined to think that Fisher's view in his later writings favors modifications of the principle of direct inference (as formulated in this book) and confirmational conditionalization, so that direct inference and credal coherence unaugmented by other principles of inductive logic would suffice for the derivation of fiducial posteriors without violating any condition on rational probability judgment.

In any case, a project of this sort has been undertaken and defended by H. E. Kyburg and is interesting in its own right. I shall discuss Kyburg's theory in chapter 16, together with A. P. Dempster's alternative proposal for abandoning confirmational conditionalization.

Although fiducial inferences have often seemed disreputable to writers on probability, induction, and statistical inference,

they have been honored by efforts on the part of serious authors to tame them by assimilating them into more orthodox views. Thus, J. Neyman initially thought his theory of confidence-interval estimation was an extension of Fisher's account of fiducial inference.³ Since in a limited domain there is a formal similarity between confidence-interval estimation and fiducial inference, one might attempt to rationalize Fisher's theory within the framework of the Neyman-Pearson theory. The Neyman-Pearson approach will be considered in chapter 17 in a general way, but this particular application of that approach will not be developed.

From H. Jeffreys onward, sympathizers with the Bayesian tradition have sought to interpret fiducial inferences in a manner consistent with the requirements of strict Bayesian doctrine.⁴ Doing so requires reconstructing such inferences so that the fiducial posteriors are also derivable via Bayes' theorem from appropriate numerically definite priors.

Following this course requires interpreting Fisher counter to the only clear intention manifested in his discussion of the fiducial argument. Nonetheless, the project is worth considering.

If a system of circumstances can be identified relative to which fiducial inferences can be prescribed in a manner consistent with the requirements of objectivist inductive logic, we may be in a position to formulate an additional principle of inductive logic to those contained in objectivist inductive logic.

To be sure, we would expect more than a consistent extension of objectivist inductive logic. The new principle would also have to be cogent.

In this chapter, fiducial inference will be "tamed" by introducing a principle of inductive logic which mandates its use in certain kinds of situations and which has been alleged to be consistent.

In chapter 15, I. Hacking's method for deriving this principle from a *prima facie* more compelling principle will be explored and the question of the consistency of fiducial inference and Hacking's theory will be investigated.

14.2 Three Cases

In contexts where fiducial arguments are to be applied, agent X knows that some chance setup is capable of responding on

a trial of kind S in one of several ways each representable by a point or a set of points in a sample space. X knows that the chance distribution over the points in the sample space is characterized by exactly one $h_\theta \in U$ where the members of U are parameterized by $\theta \in \Theta$. But he does not know which of these distributions is the correct one.

With only slight complication, it is possible to consider situations where the sample space varies for different values of θ . I shall avoid the complication and discuss only cases where the sample space remains the same relative to all values of θ .

Fisher introduced fiducial arguments for application in situations where the chance distributions are known to be continuous. I shall restrict attention to the one-dimensional case where the points in the sample space are representable by real numbers in some interval (finite or infinite) on the real line. Fisher also considered situations where points in the sample space are representable by ordered n -tuples of real numbers which are coordinates for points in a region (finite or infinite) in n -dimensional space; but, for the sake of simplicity, I shall focus only on those cases which can be transformed into one-dimensional cases.

I shall modify the previous notation so that $F_\theta(x)$ is the (statistical or chance) *cumulative distribution function* specifying the chance according to h_θ that the result of a trial of kind S is representable by a point $x' \leq x$. I take $F_\theta(x)$ to be continuous and differentiable in x (over the entire interval of points representing the sample space) so that the *density function* $f_\theta(x) = dF_\theta(x)/dx$ is defined.

I shall refer to situations of this sort as continuous cases in the subsequent discussion even though I am dealing only with one dimension.

Thanks to Hacking, examples of fiducial inferences applicable in cases where there are only a finite number of points in the sample space have been constructed,⁵ and it is also possible to construct them for situations where there are a countable infinity of points as well. Such examples are of little interest from the point of view of serious applications. However, it is helpful to begin a discussion of fiducial inference by attending to these cases. The reason is that the bare bones of the structure of such inference and the assumptions required

to employ it are thereby revealed, uncomplicated by the special features required to deal with the continuous case.

When dealing with such discrete cases, I shall let $F_\theta(x)$ be the cumulative distribution function as in the continuous case. $f_\theta(x)$, however, shall be used to represent the chance according to h_θ that the event R_x occurs on a trial of kind S . In the continuous case, of course, $f_\theta(x)$ does not represent this chance (the chance is 0). However, $f_\theta(x)dx$ approximates the chance that an event indexed by a value between x and $x + dx$ occurs.

The discrete case may be divided into two cases: the finite and the countably infinite cases. I shall consider them separately. Thus, there will be three cases to discuss: the finite, countably infinite, and continuous.

These do not exhaust all cases where fiducial arguments might be applied. However, consideration of these three will suffice for the purposes of this discussion.

14.3 The Finite Case

Fiducial inferences of the three types under consideration can be analyzed as involving three steps: a *pivotal step*, an *inversion step*, and a *commitment to irrelevance*.

In order to take the first two steps, no new principle of inductive logic need be invoked. However, the initial corpus of knowledge prior to finding out via observation which value of x occurred (i.e., for which x an event of kind R_x occurred) has to satisfy certain conditions for the pivotal step to be legitimate; it has to satisfy yet more conditions for the inversion step to be legitimate; and, in the continuous case but not the others, it must satisfy still another condition in order for the commitment to irrelevance to be made.

In case all these conditions are satisfied by the initial corpus K , the principles of direct inference, coherence, and confirmational conditionalization suffice to legitimate the pivotal step and the inversion step. To obtain a legitimate commitment to irrelevance, either X must transgress the bounds of necessitarianism or he must adopt a new principle of inductive logic which mandates adoption of the commitment to irrelevance when the corpus K satisfies the conditions that have been specified for the pivotal, inversion, and (in the continuous case) irrelevance steps.

This last strategy is the one to be considered.

Hacking's example is eminently suited to illustrate these three steps and the knowledge required to permit them to be taken in the finite case.

X knows that coin a has either a .4 chance of landing heads on a toss (and a .6 chance of landing tails) or a .6 chance of landing heads on a toss (and a .4 chance of landing tails). The parameter space contains two points $\theta = .4$ and $\theta = .6$, representing $h_{.4}$ and $h_{.6}$ respectively. The sample space contains two points: $x = 0$ for tails up and $x = 1$ for heads up.

In this example, a result of a toss of coin a is a winner if and only if either it lands heads up when $h_{.6}$ is true or lands tails up when $h_{.4}$ is true. Otherwise the result is a loser.

X knows by deductive closure that on a toss of coin a either a winner or a loser will occur.

Moreover, he knows that a winner will occur if and only if coin a lands heads up if $h_{.6}$ is true and a winner will occur if and only if coin a lands tails up if $h_{.4}$ is true. Similar remarks apply for losers.

Finally, he knows that the chance of obtaining a winner is .6 whether $h_{.6}$ is true or $h_{.4}$ is true and that the chance of a loser is, therefore, .4.

By direct inference, therefore, X is in a position to assign a degree of credence of .6 to the hypothesis that coin a comes up a winner provided he knows the coin to be tossed and that any extra information he has about the toss is known by him to be stochastically irrelevant. Assuming his initial corpus K has the information about the toss, the assignment of such degrees of credence via direct inference is the pivotal step.

The key feature of the pivotal step is the construction of a new sample space consisting of points v such that v is uniquely determined for each θ and x , and, hence, such that $v = v(x, \theta)$. This v -function is called the *pivotal function*.

In our example, the pivotal function assigns the value "winner" to (0, .4) and (1, .6) and "loser" to (1, .4) and (0, .6). For convenience, let the numerical index for winners be 1 and for losers be 0.

Of course, simply constructing such a function is not sufficient to obtain a function permitting the pivotal step. X must have the following knowledge:

(a) For every v, x , and θ , if $v(x, \theta) = v$ and h_θ is true, a result of kind T_v occurs on a trial of kind S if and only if a result of kind R_x occurs on a trial of kind S on setup a .

The chance of a result of kind T_v occurring on a trial of kind S equals the chance of a result of kind R_x occurring on a trial of kind S on setup a —i.e., $f_\theta(x) = g_\theta(v)$.

(b) $g_\theta(v) = g_{\theta'}(v) = g(v)$ for every θ and θ' .

In our example, X does have the knowledge required by (a) and (b).

By the principle of direct inference (applicable because X presumably has the requisite knowledge of stochastic irrelevance concerning trials of kind S), the degree of credence assigned to the hypothesis d_v asserting that the result is of kind T_v should be equal to $g(v)$. That is to say, for every permissible Q -function relative to K according to the credal state, $Q(d_v) = g(v)$. Thus, the pivotal step is taken.

Conditions (a) and (b) specify knowledge about the sample space generated by the function $v(x, \theta)$ necessary and sufficient for counting the function a pivotal function. Given that K contains information that a trial of kind S has occurred and that all extra information is known to be stochastically irrelevant, these conditions are necessary and sufficient to secure applicability of the pivotal step.

These remarks apply, it should be remembered, to the finite case.

The statements (a) and (b) imply that

(c) The number of values for v is the same as the number of values for x . For every θ , the function $v(x, \theta) = v_\theta(x)$ is a one-to-one mapping of the values of x onto the values of v .

Returning to the example and turning to the *inversion step*, let d_1 assert that a winner occurs on the trial of kind S and d_0 that a loser occurs. e_1 asserts that the coin lands heads up and e_0 that it lands tails up. The corpus K and e_1 entail the equivalence of $h_{.6}$ and d_1 and the equivalence of $h_{.4}$ and d_0 . K and e_0 entail the equivalence of $h_{.6}$ and d_0 and of $h_{.4}$ and d_1 .

Because the corpus K has the property just specified, credal coherence requires that

$$\begin{aligned} Q(h_{.6}; e_1) &= Q(d_1; e_1) & Q(h_{.4}; e_1) &= Q(d_0; e_1) \\ Q(h_{.6}; e_0) &= Q(d_0; e_0) & Q(h_{.4}; e_0) &= Q(d_1; e_0), \end{aligned}$$

for every permissible Q -function in $C(K)$.

Imposing these conditions on the Q -functions in the credal state is the inversion step.

As in the pivotal step, the crucial consideration is whether the corpus K satisfies conditions which justify appealing to the principles of inductive logic to mandate the inversion step.

If the corpus does satisfy these conditions, and the experiment is conducted, confirmational conditionalization implies that the unconditional Q -value for $h_{.6}$ relative to the new corpus in case e_1 is added will equal $Q(h_{.6}; e_1)$ ($= Q(d_1; e_1)$), according to the old credal state. Similar remarks apply *mutatis mutandis* when observations generate the admission of e_0 into the corpus.

The condition on K which is necessary and sufficient to justify the inversion step (given that K satisfies (a), (b), and (c) and that all the conditions on confirmational commitments adopted before are in place) for situations where the range of values for x is finite may be formulated as:

(d) For every x , $v_x(\theta) = v(x, \theta)$ is a one-to-one mapping of the values of θ onto the values of v .

I shall call this the *condition of invertibility*. Notice that this condition implies, given the conditions already imposed on K , that K and e_x entail the equivalence of h_θ and d_v for x, v , and θ such that $v(x, \theta) = v$. It also implies that the number of possible values for θ is the same as the number of values of v which, by condition (c), is the same as the number of values for x .

With these implications in force, for every x, v , and θ such that $v(x, \theta) = v$, we have that $Q(d_v; e_x) = Q(h_\theta; e_x)$ for every permissible Q -function in $C(K)$ by virtue of credal coherence.

Finally, if K_{e_x} is the deductive closure of K and e_x , then every permissible Q -function in $C(K_{e_x})$ must be such that $Q_{e_x}(d_v) = Q_{e_x}(h_\theta)$.

Returning once more to the example, it is known prior to finding out the result of the toss that the chance of a winner is .6 and, hence, by direct inference, $Q(d_1) = .6$ for every permissible Q -function in $C(K)$. Suppose X finds out that the

coin lands heads, so that e_1 is in the expanded corpus. For every permissible Q -function in $C(K_{e_1})$, $Q_{e_1}(d_1) = Q_{e_1}(h_6)$ and for every permissible Q -function in $C(K)$ $Q(d_1; e_1) = Q(h_6; e_1)$.

If we were to assume that, relative to K , the information that e_1 holds is confirmationally irrelevant to whether either a winner or a loser occurs, $Q(d_1; e_1)$ would be equated with $Q(d_1) = .6$.

Notice, however, that confirmational irrelevance can be adopted only if the principles of objectivist inductive logic and the initial corpus K do not already preclude this commitment. Suppose, for example, X knows that coin a is selected at random from an urn with 90% coins with a .4 chance of heads and 10% with a .6 chance of heads. In that case, $Q(d_1)$ would still equal .6; but $Q(d_1; e_1)$ would be less than .15.

To take care of this eventuality, another condition should be imposed on the corpus K :

(e) Relative to K , the principles of objectivist inductive logic do not prohibit Q -functions such that $Q(d_v; e_x) = Q(d_v)$ for $v = v(x, \theta)$.

If condition (e) is satisfied, v is an *irrelevance-allowing* pivotal function.

Given a corpus meeting the conditions guaranteeing that v is an invertible and irrelevance-allowing pivotal, objectivist inductive logic does not mandate assuming confirmational irrelevance in the manner illustrated by my example. But it does not prohibit it either.

However, if I were to introduce another principle of inductive logic mandating the adoption of confirmational irrelevance without further qualification, I would get into trouble.

The coin example can be used to illustrate the problem.

Suppose X knows in his initial corpus that coin a was tossed once before and had landed heads. Let K^* be the corpus such that K is the deductive closure of K^* and the information that the coin landed heads on this earlier toss.

Relative to K^* , a smoothly invertible and irrelevance-allowing pivotal function could have been constructed. If the new principle of inductive logic mandates the irrelevance step for that case, relative to K , the Q -value for h_6 would be equal to the Q -value for obtaining a winner on that first toss which is .6. Hence, the prior Q -value for h_6 for the purpose of consid-

ering the second toss is .6. Bayes' theorem then implies that $Q(h_6; e_1) = .36/.52 = 9/13$.

Observe, however, that winning and losing determine a smoothly invertible and irrelevance-allowing pivotal function relative to K for the second toss. If inductive logic mandates that, relative to K , $Q(d_1; e_1) = Q(d_1) = .6$, then $Q(h_6; e_1) = .6$ —rather than $9/13$.

Thus, if one utilizes the principle of inductive logic with regard to the first toss and corpus K^* , one cannot consistently use it for the second toss relative to K (and conversely). Yet, the contemplated principle of inductive logic mandates its own application in both cases because the pivotal functions for the first and second tosses are both smoothly invertible and irrelevance-allowing.

The source of the difficulty can be elucidated as follows:

Let f_1 assert that the coin lands heads up on the first toss and f_0 that it lands tails up.

Relative to K^* , direct inference mandates that every permissible Q^* -function satisfy

$$\begin{aligned} \text{(i)} \quad & Q^*(f_1; h_6) = .6 \\ & Q^*(f_1 \& e_1; h_6) = .36 \\ & Q^*(f_1; h_4) = .4 \\ & Q^*(f_1 \& e_1; h_4) = .16; \end{aligned}$$

$$\begin{aligned} \text{(ii)} \quad & Q^*(h_6; f_1) = \frac{.6Q^*(h_6)}{.6Q^*(h_6) + .4Q^*(h_4)} \\ & Q^*(h_6; f_1 \& e_1) = \frac{.36Q^*(h_6)}{.36Q^*(h_6) + .16Q^*(h_4)}. \end{aligned}$$

Hence, unless $Q^*(h_6)$ equals 0 or 1,

$$\text{(iii)} \quad Q^*(h_6; f_1 \& e_1) \neq Q^*(h_6; f_1).$$

Similarly,

$$\text{(iv)} \quad Q^*(h_6; e_1) \neq Q^*(h_6; f_1 \& e_1).$$

Consequently, if we assume confirmational conditionalization applies (as we are doing), it makes quite a difference whether f_1 is in K , so that K^* is a contraction of K , or is not, so that $K^* = K$, if we are attending to the Q -value assigned to h_6 conditional on e_1 . But the principle proposed as a principle of inductive logic is insensitive to this difference, and hence leads to the contradiction.

To avoid such a difficulty, a further constraint must be imposed on the corpus K :

(f) Let K^* be the weakest contraction of K satisfying conditions (a)–(d) relative to the given sample space and set U of simple statistical hypotheses. Let K be the deductive closure of K^* and f . The principles of objectivist inductive logic do not prohibit Q^* -functions relative to K^* from satisfying the condition

$$Q^*(h_\theta; e_x \& f) = Q^*(h_\theta; e_x)$$

for every θ and x .

In the coin example, if X does have in his corpus K information about previous tosses of coin a , condition (f) is violated by the corpus K .

With condition (f) in place, I am now in a position to propose a new principle of inductive logic.

The Principle of Fiducial Inference for the Finite Case: If K entails that at least and at most one of a set U of hypotheses h_θ is true and each h_θ is consistent with K , if each h_θ is a simple statistical hypothesis specifying a chance distribution over the sample finite sample space Ω on trials of kind S , if K entails that a trial of kind $S \& T$ occurs where the information that the trial is of kind T is stochastically irrelevant to hypotheses as to which event represented by a point in Ω occurs, and if K satisfies conditions (a)–(f), then

$$Q(d_v; e_x) = Q(d_v) = g(v)$$

for every Q -function in $C(K)$ and v, x, θ such that $v = v(x, \theta)$.

The multiplication theorem requires that

$$Q(h_\theta; e_x) = \frac{Q(e_x; h_\theta)Q(h_\theta)}{Q(e_x)}$$

for every Q -function in $C(K)$, provided that $Q(e_x) > 0$.

By the principle of fiducial inference, $Q(h_\theta; e_x) = g(v)$ if K satisfies the conditions of the principle. But if K satisfies those conditions, then $Q(e_x; h_\theta) = g(v) = Q(h_\theta; e_x)$. Hence, $Q(h_\theta) = Q(e_x)$ for every θ and x .

This last result obtains if and only if every h_θ is assigned an equal unconditional (i.e., prior) Q -value.

Thus, in the finite case, the principle of fiducial inference mandates a uniform prior over the rival simple statistical hypotheses.

In this respect, the principle of fiducial inference looks suspiciously like the principle of insufficient reason. However, it is not liable to the objection for the finite case which a principle of insufficient reason can face. If U contains three or more elements, the principle of insufficient reason cannot explain why it should be applied to elements of that partition rather than to some partition obtained by coarsening U .

The principle of fiducial inference in the finite case stipulates that X 's knowledge and credal state pertaining to U satisfy certain stringent conditions. Thus, even for finite U , the principle of fiducial inference in the finite case is more stringent in its domain of applicability than a principle of insufficient reason would be.

Of course, this circumstance only shows that the principle of fiducial inference in the finite case avoids obvious inconsistency. But it is possible to fix up the principle of insufficient reason in the finite case in many ways so as to avoid patent inconsistency. Doing so only reveals formal ingenuity. It does not, in any way, argue for the cogency of supplementing objectivist inductive logic by an additional principle of the sort contemplated.

At this point, however, I am not concerned to argue pro or con the adequacy of the principle but only to explain it.

However, before turning to the other cases, it should be emphasized once more that the resulting reconstruction is not in any obvious sense Fisher's. Fisher thought that fiducial inference entailed no commitment to a prior which is numerically precise. The principle of fiducial inference as formulated here does entail such a commitment. This deviation from Fisher's view is not surprising, given our quest for a method of taming fiducial inference by rendering it compatible with the other principles already imposed on confirmational commitments.

14.4 The Countably Infinite Case

The principle of fiducial inference for finite cases can be extended to apply to countably infinite cases as well by letting the sample space contain a countably infinite set of points. Otherwise the conditions are the same as in section 14.3, and

the fiducial argument can be obtained by constructing an invertible pivotal and applying the three steps.

Consider, e.g., a situation where the values of θ are all negative and positive integers and 0, as are the values of x . (Incidentally, because of the invertibility condition, when the number of values of x is countably infinite so is the number of values of θ .) Let

$$f_{\theta}(x) = \begin{cases} .5 & \text{for } \theta = x \\ (.5)^{|\theta-x|+2} & \text{for } \theta \neq x. \end{cases}$$

$v = \theta - x$ is an invertible pivotal function with $g(v) = .5$ if $v = 0$ and $(.5)^{|v|+2}$ otherwise.

By reasoning exactly analogous to that used in the previous section, the prior Q -distribution over the values of θ (and also over the values of x) must be uniform. Since there is a countable infinity of such values, they must all equal 0. Hence, in this kind of case, fiducial inference can proceed consistently only at the expense of countable additivity. Those who insist on imposing countable additivity as a condition on Q -distributions cannot consistently adopt the principle of fiducial inference for countably infinite cases. (See sections 5.11 and 12.14-12.16 for related topics.)

14.5 The Continuous Case

The fiducial argument in the continuous case has three steps, as in the two discrete cases just considered; and the legitimacy of these steps depends not only on inductive logic but on X 's corpus. However, the requirements on the corpus will be strengthened in certain ways appropriate to the problems posed by continuity.

Thus, we seek a pivotal function $v(x, \theta)$ which not only satisfies condition (a) of section 14.3, but also the following stronger requirement:

(a*) For every v', x', v'' , and θ , where $v' = v(x', \theta)$ and $v'' = v(x'', \theta)$, and when h_{θ} is true, a result indexed by a point in the interval from v' to v'' occurs on a trial of kind S if and only if a result indexed by a point in the interval from x' to x'' occurs on a trial of kind S . The chance of a result of the first kind occurring on a trial of kind S equals the chance of a result of the second occurring on a trial of kind S . That is to say,

$$|F_{\theta}(x'') - F_{\theta}(x')| = |G_{\theta}(v'') - G_{\theta}(v')|.$$

Condition (b) is strengthened as follows:

(b*) $G_{\theta}(v) = G_{\theta'}(v)$ for every θ and θ' .

(a*) and (b*) presuppose that $v(x, \theta)$ satisfies

(c*) The number of values for v is the same as the number of values for x . $v(x, \theta) = v_{\theta}(x)$ is a strictly monotonic differentiable function of x with continuous derivative which maps onto the values of v (and, hence, is one-to-one onto v). The density $g_{\theta}(v)$ exists and $g_{\theta}(v)dv = f_{\theta}(x)dx$. For every θ and θ' , $g_{\theta}(v) = g_{\theta'}(v) = g(v)$.

Conditions (a*) and (b*) (or condition (c*)) specify conditions for a pivotal function in the continuous case.

In a similar vein, condition (d) is strengthened as follows:

(d*) For every x , $v_x(\theta) = v(x, \theta)$ is a strictly monotonic differentiable function of θ with a continuous derivative that maps all values of θ onto the range of values for v .

A pivotal function satisfying this condition is *smoothly invertible* (following J. Tukey).⁶

If the pivotal function is smoothly invertible, the values of θ fall in some interval on the real line. For $v' = v(x, \theta')$ and $v'' = v(x, \theta'')$, K and e_x entail that the value of v falls in the interval from v' to v'' if and only if the true value of θ lies between θ' and θ'' . Let $H_x(\theta)$ be the cumulative distribution function for θ representing the (or a) Q -distribution for θ conditional on e_x . Smooth invertibility implies that either

$$H_x(\theta) = G_x(v)$$

or

$$H_x(\theta) = G_x(v).$$

If $G_x(v)$ is continuous and differentiable with respect to v with continuous derivative, then $H_x(\theta)$ is continuous and differentiable with respect to θ with continuous derivative. In that case, $h_x(\theta)d\theta = g_x(v)dv$.

To license the step at which confirmational irrelevance is adopted, the smoothly invertible pivotal function v must also be irrelevance-allowing. In the continuous case, this condition should be strengthened as follows:

(e*) Relative to K , the principles of objectivist inductive logic do not prohibit Q -functions where the distribution over values of v are representable by a cumulative distribution function $G_x(v) = G(v)$.

Because $G(v)$ is continuous and differentiable in v with continuous derivative $g(v)$, so is $G_x(v)$. Thus, $g_x(v) = g(v)$. Hence, $H_x(\theta)$ is continuous and differentiable in θ with continuous derivative $h_x(\theta)$.

By smooth invertibility, we have

$$(1) \quad g_x(v)dv = h_x(\theta)d\theta.$$

By the conditions on pivotal functions,

$$(2) \quad g(v)dv = f_\theta(x)dx.$$

By the adoption of confirmational irrelevance,

$$(3) \quad g_x(v)dv = g(v)dv.$$

From the properties already derived for smoothly invertible pivotal functions, it follows that, for fixed v , $\theta_v(x) = \theta$ is a strictly monotonic differentiable function of x with continuous derivative where $v(x, \theta) = v$ if and only if $\theta_v(x) = \theta$. If $f_v(x)$ is the density for values of x conditional on d_v and $h_v(\theta)$ the corresponding density for θ ,

$$(4) \quad f_v(x)dx = h_v(\theta)d\theta.$$

By the multiplication theorem,

$$(5) \quad f_v(x)g(v) = g_x(v)f(x),$$

$$(6) \quad f_v(x) = f(x),$$

because $g_x(v) = g(v)$ by (3).

By similar reasoning,

$$(7) \quad h_v(\theta) = h(\theta).$$

From (4) and (6) and (7), it follows that

$$(8) \quad \frac{\partial \theta}{\partial x} = \frac{f_v(x)/h_v(\theta)}{f(x)/h(\theta)}.$$

Condition (8) is derived from the assumption of confirmational irrelevance. But it asserts that the function $v(x, \theta) = v$ must be such that $\partial \theta / \partial x$ is a product of a function of θ alone

and of x alone. Hence, if the pivotal is irrelevance-allowing in the sense of (e*), this condition must be satisfied even if confirmational irrelevance is not adopted.

But $\partial \theta / \partial x$ (where θ is taken to be a function of x and v) is the product of a function of x alone and of θ alone if and only if

$$(9) \quad \theta(x, v) = J(B(x) + C(v)).$$

This in turn holds if and only if

$$(10) \quad v(x, \theta) = I(D(x) + E(\theta)).$$

Hence, a smoothly invertible pivotal is irrelevance-allowing if and only if (10) is satisfied.

If condition (10) holds, one can find one-to-one transformations $w(v)$, $y(x)$, and $\phi(\theta)$ such that $w(y, \phi) = y + \phi$. w is a smoothly invertible pivotal trivially satisfying (10).

One further condition must be imposed on K . A condition (f*) analogous to (f) for the discrete cases must be added. The only modification needed is that the Q^* -function used in (f) should be replaced by appropriate cumulative distribution functions for θ conditional on e_x & f and e_x , respectively.

I am now in a position to formulate a principle of fiducial inference for the continuous case.

The Principle of Fiducial Inference for the Continuous Case: If K entails that at least and at most one of a set U of hypotheses h_θ is true and each h_θ is consistent with K , if each h_θ is a simple statistical hypothesis specifying a chance distribution over a sample space consisting of all points in some interval Ω of the real line on trials of kind S , if K entails that a trial of kind S & T occurs where the information that the trial is of kind T is stochastically irrelevant to hypotheses as to which event represented by a point in Ω occurs, and if K satisfies (a*)-(f*), then $g_x(v) = g(v)$ for every Q -function in $C(K)$ and v, x, θ such that $v = v(x, \theta)$.

This principle may be illustrated by situations where it is known that x is normally distributed with unit population variance and unknown population mean θ , i.e., where

$$f_\theta(x) = \frac{1}{\sqrt{2\pi}} e^{-(x-\theta)^2/2}.$$

Let $v = x - \theta$. Then, for every θ ,

$$g_{\theta}(v) = g(v) = \frac{1}{\sqrt{2\pi}} e^{-v^2/2}.$$

The pivotal is smoothly invertible and irrelevance-allowing. Given evidence that on a single trial a result of kind R_x has occurred, the posterior distribution for θ is given by

$$h_x(\theta) = \frac{1}{\sqrt{2\pi}} e^{-(x-\theta)^2/2}.$$

In this example, by letting $\phi = -\theta$, we obtain $v = x + \phi$. And if we have any irrelevance-allowing, smoothly invertible pivotal, we can obtain $w = y + \phi$. For the sake of our notation, let $x = y$, $w = v$, and $\theta = \phi$, so that $v = x + \theta$. It then follows that

$$(11) \quad \frac{\partial v}{\partial x} = \frac{\partial v}{\partial \theta} = 1 \quad \frac{\partial \theta}{\partial x} = 1.$$

From (11) and (8) we derive

$$(12) \quad \frac{h(\theta)}{f(x)} = 1$$

for every θ and x .

This can happen if and only if the densities for θ and x are everywhere uniform. If the initial corpus K is such that the principle of fiducial inference in the continuous case is applicable, then one can always find a transformation of pivotal, random variable, and parameter such that the prior distribution of the parameter is uniform.†

In our example, $-\infty < x < \infty$. Hence, $-\infty < \theta < \infty$. Hence, as in the countably infinite case, countable additivity must be violated.

† The results obtained here are based on D. V. Lindley's paper "Fiducial Distributions and Bayes' Theorem," *J. Royal Stat. Soc.*, ser. B, v. 20 (1958), pp. 102-107. Lindley considers the continuous case and claims to specify necessary and sufficient conditions for the consistency of the fiducial argument in the one-parameter case. Lindley's claim is mistaken, as T. Seidenfeld has shown and as will be explained in the next chapter. What is true is that no inconsistency will appear with Bayesian principles as long as all "data" are outcomes of repetitions of trials of a single kind S . I have exploited Lindley's conditions in formulating the principle of fiducial inference for the continuous case.

Suppose the kind of trial yielding the normally distributed outcomes is repeated n times so that the data consist of n values of x . A point in the sample space is an n -tuple in an n -dimensional space. However, the information conveyed by any such point can be condensed into a specification of the value of the sample mean \bar{x} and the sample variance s^2 . The chance according to h_{θ} of obtaining a given value \bar{x} for the sample mean and a given value s^2 for the sample variance is equal to $k(\bar{x}, s^2)$ times the chance of obtaining an n -tuple entailing such a mean and sample variance where $k(\bar{x}, s^2)$ is constant for all values of θ . Thus, given the information about the value of \bar{x} and s^2 , the extra information as to which n -tuple of values of x has occurred from among those entailing these values is confirmationally irrelevant to the value of θ in the strong sense. Hence, we may ignore this extra information and focus on the value of \bar{x} and s^2 .

However, the chance of \bar{x} is independent of the chance of s^2 on a series of n repetitions. Hence, if we consider the kind of trial which is a series of n repetitions of the trial of kind S in which the sample variance is s^2 , the extra information that the sample variance is s^2 is known to be stochastically irrelevant.

If we begin with a corpus K where it is known that series of n repetitions of trials of kind S yielding a sample variance s^2 has taken place, we can let $v = \bar{x} - \theta$; v will be a smoothly invertible irrelevance-allowing pivotal function, and a fiducial argument can be developed in accordance with our principles.

As before, the prior for θ will be uniform. Indeed, if inconsistency is not to threaten, this must be so.

This account of fiducial inference in the continuous case is far too weak to handle all the problems Fisher wished to consider. For example, the fiducial distribution for the normal mean when both mean and variance are unknown involves a use of Student's t -statistic. The fiducial argument which applies in this case requires stronger principles than I have proposed here.

A more extensive treatment of fiducial arguments has been given by T. Seidenfeld.⁷ The philosophical issues to be raised here, however, do not require our entering into the technicalities involved in a more general treatment. The more technical

issues are philosophically interesting; but I mean to avoid them here.

14.6
Is Tame
Fiducial
Inference
Convincing?

Tame fiducial inference must perforce yield numerically definite posteriors that presuppose numerically definite priors. Fisher did not wish to understand fiducial inference in this manner. Hence, we should not suppose that his fiducial inferences were tame ones.

But whether Fisher would have adopted the principle of fiducial inference as a principle of inductive logic or not, it is entirely reasonable to inquire into its adequacy.

In the next chapter, an attempt to rationalize adoption of a principle of fiducial inference by an appeal to an analysis of the concept of chance will be considered. If such an attempt could be rendered acceptable, the principle of fiducial inference would be on much the same footing as the principle of direct inference. I do not believe such a rationalization can be achieved, but the attempt is worth exploring.

15.1
Likelihood and
Evidential
Support

According to R. A. Fisher, apart from the simple test of significance, "there are to be recognized and distinguished, between the levels of certain knowledge and total nescience, two well-defined levels of logical status for parameters lying on a continuum of possible values, namely that in which the probability is known for the parameter to lie between any assigned values, and that in which no probability statements being possible, or only statements of inequality, the Mathematical Likelihood of all possible values can be determined from the body of observations available."¹

In this passage and Fisher's various elaborations of it, he fails to clarify his intent in speaking of the "logical status" for the parameters in question. He does seem to think, however, that when numerically precise probability judgments cannot be made concerning rival statistical hypotheses represented by values of a continuous parameter, information about likelihoods may, nonetheless, be usefully exploited.

From a strict Bayesian point of view, there is no doubt that likelihood has importance; for posterior credal probabilities are a function of prior probabilities and likelihoods. Likelihood may be construed as furnishing a measure of the contribution of the data of experimentation to the modification of the credal state from the prior to the posterior one. Because credal probabilities have a definite use in inquiry and deliberation, likelihoods acquire a derivative import as well.

One does not have to endorse strict Bayesianism to acknowledge this import. Provided that the prior credal state is not maximally indeterminate, likelihoods can be used to determine modifications of that state due to the contribution of the data of experimentation.

Trouble arises, however, when the prior credal state is maximally indeterminate. In that case, the data make no contribution to the determination of the posterior credal state except

through deductive logic. Likelihood cannot, therefore, make any contribution.

Yet, Fisher does seem to hold that in such cases likelihoods supply relevant and useful information. What can be meant by this?

It is not helpful to claim that likelihood provides a measure of evidential support, relevance, or confirmation. All of these terms have been used in so many ways and with such obscurity that unless some indication is given of how such assessments of support are to be employed in inquiry and deliberation, this characterization of the import of likelihood is hopeless.†

The difficulty is not that measures of likelihood fail to provide probabilistic appraisals of hypotheses. No doubt many important modes of appraisal are not probabilistic.

Thus when seeking to expand his corpus, an investigator should be concerned to evaluate rival potential answers with respect to how worthy they are of being added. In at least one sense, such appraisal is an evaluation of support, confirmation, or the like. Moreover, such evaluation is important. Nonetheless, it is not probabilistic. The support for a hypothesis in this sense is, according to the approach I favor, represented by the permissible values for expected epistemic utility allocated to the rival potential answers.² Such support depends on judgments of credal probability, but is a mode of appraisal which itself is distinctly nonprobabilistic.

There are other modes of appraisal useful for other purposes which are also nonprobabilistic, such as, e.g., G. L. S. Shackle's measures of potential surprise. These too may be called measures of support.

The point is that calling a mode of appraisal a way of measuring or assessing support is not helpful. A specification of the use or function of the mode of appraisal in inquiry and deliberation is needed, whether the method is probabilistic or not.

The observation applies to likelihood as to other measures. Granted that likelihoods function in a clearly identifiable and important way in deriving posteriors from priors when priors

† Thus, in spite of its elegance, sophistication, and comprehensiveness, A. W. F. Edwards' book *Likelihood* (Cambridge: Cambridge University Press, 1972) fails to furnish useful clues as to how likelihood appraisals are to be used in inquiry and deliberation. We are told only that they furnish evaluations of support or confirmation.

are definite, what purpose do they serve when priors are maximally indeterminate? Fisher seems to maintain that they serve an important purpose in such contexts. But it is simply insufficient to say they measure support in such cases.

The position might be taken that conclusions reached in science are evaluated according to different standards than decisions taken in practical deliberation. In part I agree. The aims of scientific inquiry differ from the aims of moral, economic, political, and other forms of practical deliberation. Yet, there are certain general criteria for rational choice that apply to decision making in scientific inquiry and practical deliberation alike. Unless apologists for likelihood as an index of support can show how likelihood is relevant according to these general criteria or is relevant to the special aims of science, their views remain incomplete.

The only clear role for likelihoods identified thus far has been in the determination of posterior probabilities from priors on the data of experimentation via Bayes' theorem and confirmational conditionalization.

Likelihoods make no contribution to the modification of priors on the basis of data if the prior credal state is maximally indeterminate and objectivist inductive logic is endorsed. Likelihoodists face a predicament analogous to the one confronting fiducialists: Either new principles need to be added to objectivist inductive logic sufficient to render prior credal states less than maximally indeterminate on the basis of considerations of inductive logic and the available evidence, or some of the conditions already imposed on confirmational commitments (e.g., confirmational conditionalization and direct inference) have to be modified.

I. Hacking has sought to make likelihood central to an account of evidential support and a reconstruction of fiducial inference. As I understand him, his approach seeks to take fiducial inference and likelihoodism in a Bayesian manner; so he should be regarded as following the first strategy.

Hacking does not, to be sure, formulate his theory using the apparatus I have constructed. He begins by endorsing Koopman's postulates for comparative probability as a "logic of support" and adds further conditions.³ In the course of his discussion, he rather abruptly adds the stipulation that when a quantitative determination of support can be given, it should

obey Kolmogoroff's axioms for probability including countable additivity.†

Hacking fails to indicate whether, on his view, there should be at least one numerical representation obeying the requirements of the calculus of probability when support is not representable by a unique measure. I adopt the charitable view that he is committed to credal coherence and consistency. I shall also suppose (without any textual warrant) that he is committed to credal convexity and, indeed, that on his view X 's credal state relative to K should be the largest convex set of Q -distributions relative to K compatible with some quasi ordering obeying Koopman's axioms and permitting a probabilistic representation. In Hacking's discussion of fiducial inference, there is substantial indication of his at least implicit commitment to confirmational conditionalization.

These attributions are far more precise than anything that can be supported by Hacking's text. For this reason, the following discussion runs the risk of placing Hacking in a Procrustean bed. Nonetheless, the reconstruction of Hacking's theory that emerges is interesting in its own right and instructive concerning strategies for making likelihood central in inverse inference.

The resulting theory is stronger than mine, for it requires that credal states be describable completely in terms of quasi orderings and that permissible Q -functions satisfy countable additivity. The first difference will prove unimportant in this discussion. The second will play a small role later on.

Hacking is committed to a principle of direct inference (which he calls "the frequency principle").⁴ But he is not objectivist. He explicitly introduces two additional principles into his logic of support (i.e., his inductive logic).

Both of these additional principles appeal explicitly to likelihood. One of them, the principle of irrelevance, turns out to be unnecessary for the purposes Hacking uses it for. This will be shown later. It is important to understand at the outset,

† I. Hacking, *Logic of Statistical Inference* (Cambridge: Cambridge University Press, 1965), pp. 134–135. Hacking shies away from claiming that when support goes quantitative it should always conform to Kolmogoroff's axioms; but he does assume this to be so "for the special case of support for statistical hypotheses by statistical data" (*ibid.*, p. 134). That concession will suffice for present purposes.

however, that Hacking's new principles license an appeal to likelihood in a Bayesian taming of the fiducial argument.

Hence, when objectivist inductive logic fails to rule out any prior distribution over rival simple statistical hypotheses rendering the prior credal state maximally indeterminate if one remains an objectivist necessitarian, Hacking uses likelihood appraisals to cut down the permissible prior distributions so that likelihoods contribute to the determination of priors as well as to the derivation of posteriors from priors via Bayes' theorem.

15.2 The Law of Likelihood in the Discrete Case

The first of the two principles added to objectivist inductive logic by Hacking is his so-called *law of likelihood*.⁵ I shall first explain this principle as it applies in discrete (finite or countably infinite) cases and subsequently consider the continuous case. I shall show how principles of fiducial inference formulated in the previous chapter may be derived from Hacking's law of likelihood and the principles of objectivist inductive logic. Hacking himself appealed also to a *principle of irrelevance*.⁶ I shall discuss this principle and show that it is unnecessary for the purpose of deriving principles of fiducial inference.

As in the previous chapter, X knows that at least and at most one of a set U of hypotheses parameterized by values of θ is true. Each specifies a chance distribution over the same space Ω . Let e_x assert that a result of kind R_x occurs on a specific trial of kind S .

Let W be the set of hypotheses of the form e_x & h_θ . We shall consider only hypotheses equivalent given K to elements of the σ -algebra of propositions $S(W)$ generated by the set W .

A hypothesis $g \in S(W)$ is a *simple joint proposition* if and only if it asserts that a result of a certain kind occurs on the particular trial of kind S under consideration and specifies the chance of a result of that kind occurring on a trial of kind S or is equivalent given K to such a proposition.

All elements of W are simple joint propositions. But so are hypotheses of the form e_α & h_θ where e_α asserts that a result occurs indexed by a point in $\alpha \subseteq \Omega$. If there is a pivotal function $v = v(x, \theta)$, any hypothesis d_v is a simple joint proposition or equivalent given K to one; for the corpus K entails h that specifies the chance distribution over the values of the

pivotal function and d_v & h is, therefore, equivalent given K to d_v .

For the present, attention will be focused on cases where Ω is finite or countably infinite. To simplify discussion, I shall consider only cases where h_θ specifies a countably additive chance distribution over Ω . This is all that is considered in most contexts and, in any case, the restriction could be lifted if necessary. Moreover, I am not supposing that the Q -distributions over the hypotheses in $S(W)$ must be countably additive but only that the chance distributions according to the various simple statistical hypotheses are countably additive.

In section 4.3, I took the likelihood (or a likelihood) of h_θ on e_x in a credal state B relative to K to be any permissible value of $Q(e_x; h_\theta)$. Moreover, when K is such that all the extra information about the specific trial of kind S is known to be stochastically irrelevant, $Q(e_x; h_\theta) = f_\theta(x)$ for all permissible Q -functions where $f_\theta(x)$ is the chance of obtaining a result of kind R_x on a trial of kind S . This is implied by the principle of direct inference. Hence, $f_\theta(x)$ functions as the likelihood of h_θ on e_x in the appropriate state of knowledge.

Instead of defining likelihoods for simple hypotheses of the type h_θ belonging to U relative to data points specifying values in the sample space, Hacking defines them for the simple joint propositions of the type e_x & h_θ belonging to W or, more generally, for any simple joint proposition in $S(W)$. Indeed, the still more general characterization of *HL-likelihoods* conditional on f is as follows:

Let f be any hypothesis consistent with K and let g be any hypothesis equivalent given K and f to a simple joint hypothesis in $S(W)$. The *HL-likelihood of g conditional on f* $HL(g; f)$ is equal to the chance of an outcome of the kind asserted by g to occur according to the simple hypothesis entailed by g (given K and f).

Thus, $HL(e_x \& h_\theta; e_x) = HL(h_\theta; e_x)$ equals the chance $f_\theta(x)$ of obtaining a result of kind R_x on a trial of kind S . Hacking's conception yields the usual notion of likelihood as a special case.

However, $HL(e_x \& h_\theta; h_\theta)$ also equals the chance $f_\theta(x)$ and it is equivalent to $HL(e_x; h_\theta)$ which is the *HL-likelihood of e_x on h_θ* .

The unconditional *HL-likelihood* of g is $HL(g)$, where g is equivalent given K to some simple joint proposition in $S(W)$. $HL(g)$ is the chance of the appropriate result according to the simple statistical hypothesis asserted by g to be true: $HL(e_x \& h_\theta) = f_\theta(x)$.

Thus,

$$HL(h_\theta; e_x) = HL(e_x; h_\theta) = HL(e_x \& h_\theta) = f_\theta(x).$$

Let f be consistent with K and let V be a set of hypotheses which are simple joint propositions in $S(W)$, exclusive and exhaustive relative to K and f and each consistent with K and f . Then

(i) If V is finite and Q is permitted relative to K by the principles of objectivist inductive logic, Q agrees with the conditional *HL-function* $HL(g; f)$ where $g \in V$ if and only if for every g and g' in V , $Q(g; f) \geq Q(g'; f)$ if and only if $HL(g; f) \geq HL(g'; f)$.

(ii) If V is countably infinite and Q is permitted relative to K by the principles of objectivist inductive logic, Q agrees with the conditional *HL-function* $HL(g; f)$ where $g \in V$ if and only if for every g and g' in V , $m_{K,f}(g) \geq m_{K,f}(g')$ if and only if $HL(g; f) \geq HL(g'; f)$ where the $m_{K,f}$ -function is a σ -finite measure over V characterizing the conditional Q -distribution.

Hacking formulates his law of likelihood for discrete cases as asserting that if the *HL-likelihoods* of g and g' given f exist, f supports g better than it supports g' if the likelihood of g given f exceeds the likelihood of g' given f .⁷ I interpret his somewhat enigmatic formulations as implying that ranking of hypotheses with respect to credal probability should agree with ranking with respect to *HL-likelihood*.

When V is finite or countably infinite, this stipulation is given precise formulation by requiring conditional Q -distributions over elements of V permissible in $C(K)$ to those which agree in sense (i) or sense (ii) with the given conditional *HL-function*.

Sometimes, however, the corpus K and the principles of objectivist inductive logic might prohibit any such conditional Q -distribution from being permissible. In that event, a clash between objectivist inductive logic and the law of likelihood

emerges. But this clash is avoidable. Agreeability will be mandated only in cases where objectivist inductive logic and K permit it. This leads to the following rendition of Hacking's proposal:

Law of Likelihood for Discrete Cases: Let V be a set of simple joint propositions in $S(W)$ exclusive and exhaustive relative to K and f (where f is consistent with K) and each consistent with K and f . V is finite or countably infinite. If K and the principles of objectivist inductive logic permit at least some Q -distributions over V conditional on f to agree with the conditional HL -function $HL(g; f)$, then, for every $Q \in C(K)$, $Q(g; f)$ agrees (in sense (i) or sense (ii) depending on whether V is finite or countably infinite) with $HL(g; f)$.

Hacking formulated the law of likelihood prior to introducing his version of the principle of direct inference; this principle by itself has much of the power of the principle of direct inference.

Suppose X knows that coin a is to be tossed twice with a chance of heads on each toss equal to .5 and X also knows that the tosses are stochastically independent.

In that case, if X also knows that coin a is tossed twice, the HL -likelihood of the simple joint proposition that the coin lands heads twice is .25 as is the HL -likelihood of the hypothesis that it lands tails twice, lands heads first and tails second, and tails first and heads second. Relative to K these hypotheses form an exclusive and exhaustive partition V of simple joint propositions each with equal HL -likelihood. The law of likelihood implies that they should all bear equal Q -value and, hence, by coherence should bear Q -value of .25. This, of course, is already mandated by the principle of direct inference.

Hacking, however, endorses the frequency principle (i.e., principle of direct inference) and alleges that it "seems so universally to be accepted that it is hardly ever stated."⁸

The truly significant applications of the law of likelihood construed as an addendum to the principles of objectivist inductive logic concern inverse inference. Thus, if X is concerned with the range of Q -distributions over hypotheses h_θ in U conditional on e_x , then V consists of all simple joint propositions e_x & h_θ where x is fixed and θ is allowed to vary over

all values consistent with K . When V is finite or countably infinite (i.e., when U is finite or countably infinite),

$$Q(h_\theta; e_x) \geq Q(h_{\theta'}; e_x)$$

if and only if

$$HL(h_\theta; e_x) \geq HL(h_{\theta'}; e_x).$$

This consequence is already substantially stronger than what objectivist inductive logic mandates.

I have, however, said nothing about cases where V is larger than a countably infinite set. Suppose, for example, that X does not know the true chance p of coin a landing heads on a toss but knows that it has been tossed once and has landed heads. There is, in that case, a definite HL -likelihood determined for hypothesis h_p specifying a precise value for the chance. Moreover, the likelihood function $HL(h_p)$ is a continuous function of p . (Indeed, $HL(h_p) = p$.) Shall we say that relative to K , the higher the value of p the greater the support? Hacking's formulation of the law of likelihood seems to say this. But it is utterly unclear what this is supposed to mean. It cannot mean that the Q -value for h_p increases with p , since that value may be positive for at most a countable number of values of p . Since the HL -function is continuous in p , it may be urged that the density $f(p)$ increases continuously with the HL -likelihood.

But this proposal runs into difficulties. Let, e.g., $r = \log p$. The HL -likelihood for the hypothesis h_p is also the HL -likelihood for the hypothesis g_r , where $r = \log p$. We would have as much right to require that the density $f^*(r)$ increase continuously with the HL -likelihood. But $f^*(r) = pf(p)$. Suppose that $p > p'$. According to the first prescription, $f(p) > f(p')$. According to the second, $f^*(r) > f^*(r')$. But this holds if and only if $pf(p) > p'f(p')$. This latter result is compatible, however, with $f(p) < f(p')$.

Fortunately in the reconstruction of the fiducial argument utilizing Hacking's law of likelihood, we do not have to deal with cases of this sort. When the sample space is finite, the invertibility requirement on pivotals implies that the range of parameter values must be finite. When the sample space contains a countable infinity of points, so does the parameter

space. For this reason, it seems to be desirable to minimize controversy and restrict applicability of Hacking's law to cases where it can be clearly formulated and where it seems most compelling.

**15.3
The Law of
Likelihood and
Fiducial Infer-
ence in the
Finite Case**

Let the sample space Ω be finite, and let $v(x, \theta)$ be an invertible and irrelevance-allowing pivotal function in the sense of chapter 14. Let the corpus satisfy condition (f) of that chapter.

For fixed v , consider all simple joint propositions e_x & h_θ such that $v = v(x, \theta)$. The unconditional *HL*-likelihoods of these propositions all have the value $g(v)$. By the law of likelihood, all such joint propositions have equal *Q*-value. Furthermore, by the properties of the pivotal function, the chance of obtaining a result of kind R_x on a trial of kind S according to h_θ is equal to $f_\theta(x) = g(v)$. Hence,

$$Q(e_x \& h_\theta) = Q(e_x; h_\theta)Q(h_\theta) = g(v)Q(h_\theta),$$

by credal coherence and direct inference.

From this it follows that $Q(h_\theta)$ is a constant k for all θ . Therefore,

$$Q(e_x \& d_v) = Q(e_x \& h_\theta) = kg(v)$$

$$Q(d_v; e_x) = \frac{kg(v)}{Q(e_x)}.$$

But for fixed x and allowing v to vary,

$$\sum Q(d_v; e_x) = 1$$

where the sum is over all values of v . Hence, $k = Q(e_x) = Q(h_\theta)$ and

$$Q(d_v; e_x) = g(v)$$

$$= Q(d_v) \text{ by direct inference.}$$

Thus, the assumption of confirmational irrelevance of e_x for the d_v 's needed for the fiducial argument in the finite case is obtained using the law of likelihood for discrete cases as the sole addendum to objectivist inductive logic.

The law of likelihood is used only once in the argument—namely, to establish that the *HL*-likelihoods of all simple joint propositions of type e_x & h_θ where $v(x, \theta) = v$ for some fixed value of v are equal and, hence, bear equal *Q*-value.

**15.4
The Law of
Likelihood and
Fiducial Infer-
ence in the
Countably
Infinite Case**

Let the sample space be countably infinite and $v(x, \theta)$ an invertible and irrelevance-allowing pivotal function in the sense of chapter 14. Condition (f) of that chapter is also satisfied.

For fixed v , consider all simple joint propositions e_x & h_θ such that $v = v(x, \theta)$. The *HL*-likelihoods of these propositions all equal $g(v)$. Hence, by the law of likelihood they should also bear equal σ -finite measure, i.e., $m_{K(e_x \& h_\theta)} = c(v)$.

By direct inference and the conditions on the pivotal, $Q(e_x; h_\theta) = g(v)$ for $v = v(x, \theta)$.

By the stipulations on conditional σ -finite measures of section 5.11,

$$m_{K, h_\theta}(e_x) = Q(e_x; h_\theta) = g(v)$$

and

$$m_{K, h_\theta}(e_x) = \frac{m_K(e_x \& h_\theta)}{k(\theta)},$$

where $k(\theta)$ depends only on θ . Hence, $c(v) = g(v)k(\theta)$.

But as we have seen, if v is held constant and x and θ varied so that $v = v(x, \theta)$, then

$$m_K(e_x \& h_\theta) = c(v),$$

which is constant as long as v remains fixed. Hence, $k(\theta) = k$ for all values of θ .

Furthermore,

$$m_K(e_x \& h_\theta) = m_K(e_x \& d_v) = kg(v)$$

for $v = v(x, \theta)$.

By the stipulations on conditional σ -finite measures of section 5.11, $m_{K, e_x}(d_v)$ is directly proportional to $m_K(e_x \& d_v) = kg(v)$ as x is held fixed and v allowed to vary.

Hence, $m_{K, e_x}(d_v)$ is directly proportional to $g(v)$ —i.e., is equal to $wg(v)$ for some constant w . Since the distribution for v sums to 1 in typical cases,

$$Q(d_v; e_x) = Q(d_v) = g(v).$$

Of course, if the σ -finite values for the d_v 's conditional on e_x yielded a divergent sum, then

$$Q(d_v; e_x) = Q(d_v) = g(v) = 0.$$

This case is not seriously explored in applications of fiducial

arguments. Probability distributions over pivotal values normally are expected to obey countable additivity.

**15.5
The Principle
of Irrelevance
in the Discrete
Case**

Let V be a set of simple joint propositions in $S(W)$ exclusive and exhaustive relative to K and each consistent with K . Consider some condition f consistent with K such that the elements $g \in V$ are all consistent with K and f .

The unconditional HL -likelihood of $g \in V$ is equal to $HL(g; f)$. By the law of likelihood, $Q(g) \geq Q(g')$ for every $Q \in C(K)$, if and only if $Q(g; f) \geq Q(g'; f)$.

From this it does not follow, however, that $Q(g; f) = Q(g)$.

Hacking's *principle of irrelevance* for discrete cases mandates that when V is finite or countably infinite, $Q(g; f) = Q(g)$ for every $g \in V$.⁹

Consider once more a situation where K insures the existence of an invertible and irrelevance-allowing pivotal satisfying the condition (f) of chapter 14 in the finite case.

Let V consist of all hypotheses specifying a value of the pivotal function—i.e., of all hypotheses of the form d_v . Let the condition be e_x for some fixed x . Then $HL(d_v; e_x) = HL(h_\theta; e_x)$, where $v = v(x, \theta)$. This HL -likelihood equals $f_\theta(x) = g(v)$. But $HL(d_v) = g(v)$.

Thus, the conditions for applying the principle of irrelevance are met. We require, therefore, that $Q(d_v; e_x) = Q(d_v)$ for every $Q \in C(K)$.

This, of course, is the critical step of adopting confirmational irrelevance.

Notice that if V is countably infinite, the same result is obtainable using m_K -functions.

Hacking derives his fiducial argument in discrete cases (actually only in finite cases) from an appeal to his principle of irrelevance. It is clear that the principle of irrelevance is not itself a consequence of the addition of the law of likelihood to the principles of objectivist inductive logic.

It should also be clear that it is not necessary to introduce this principle in order to rationalize fiducial inference in discrete cases.

The matter is of some importance.

In the first place, it may turn out easier to defend the co-gency of the law of likelihood which is weaker in the relevant respects than the principle of irrelevance.

Secondly, should some difficulty appear in the consequences of using the fiducial argument, something may have to be modified. It is important to appreciate that abandoning the principle of irrelevance is of no help if the law of likelihood and objectivist inductive logic are retained.

**15.6
The Law of
Likelihood in
the Continu-
ous Case**

Let the sample space Ω be some interval of the real line where, for each θ , the chance distribution according to h_θ is representable by a continuous density $f_\theta(x)$. As before, W is the set of hypotheses of the type e_x & h_θ and $S(W)$ is a σ -algebra of propositions generated by W and all hypotheses equivalent given K to these.

Each element of W is representable by a point (x, θ) in the joint space $\Omega \times \Theta$.

$\eta(x, \theta) = f_\theta(x)$ is the HL -density assigned the point (x, θ) in the joint space. $\eta(x, \theta; f) = f_\theta(x)$ is the HL -density assigned the point (x, θ) in the subspace of all points representing simple joint propositions consistent with f . It is the conditional HL -density function.

Let $y(x)$ be any strictly monotonic differentiable function of x with continuous derivative. We may construct a new joint space of points (y, θ) to represent elements of W . $\eta^*(y, \theta)dy/dx = \eta(x, \theta)$ is the HL -function for this new space. A similar definition applies to conditional HL -functions.

Let f be some hypothesis consistent with K ; let V be a set of elements of W exclusive and exhaustive relative to K and f , each element representing a hypothesis consistent with K . Moreover, one and only one element of V is associated with a value of x in the range of values of x .

If V is as specified and Q is permitted relative to K by the principles of objectivist inductive logic, Q agrees with the conditional HL -density function $\eta(x, \theta; f)$ if and only if the HL -density function ranges continuously over the points in V , Q conditional on f is representable by a continuous density function ranging over the values of x of the form $f(x; f)$, and for every x and x' , $f(x; f) \geq f(x'; f)$ if and only if $\eta(x, \theta; f) \geq \eta(x', \theta'; f)$ where (x, θ) and (x', θ') are points in V .

Notice that the Q -distribution may also be representable by a density $f^*(y; f)$ where y is a strictly monotonic and differentiable function of x with continuous derivative and

$$f^*(y; f) = f(x; f) \frac{dx}{dy} = \eta(x, \theta; f) \frac{dx}{dy} = \eta^*(y, \theta; f).$$

Thus, ambiguity in the condition of agreement is avoided.

We are now in a position to state the law of likelihood for continuous cases.

Law of Likelihood for Continuous Cases: Let V be as stipulated. If K and the principles of objectivist inductive logic permit at least some Q -distributions over V conditional on f to agree with the conditional HL -function on f , then, for every $Q \in C(K)$, the Q -distribution over elements of V conditional on f is representable by a conditional density agreeing with the HL -density on f .

This formulation of the law of likelihood for the continuous case differs from Hacking's in several respects.¹⁰ In the first place, I have restricted my definition of HL -densities to apply to points in the joint space representing simple joint propositions in W . Moreover, the HL -density associated with such an element of W is not unique but depends on the joint-space representation of the simple joint proposition. If $y = y(x)$, then e'_y & h_θ is equivalent given K to e_x & h_θ . The point (x, θ) represents the same element of W as (y, θ) does in its different joint space. But the HL -density assigned (x, θ) differs from that assigned (y, θ) .

Hacking assigns neither HL -densities nor HL -likelihoods to the elements of W . Given any two such hypotheses, he introduces a likelihood ratio. His law of likelihood states that one joint proposition is better supported by the data if the likelihood ratio of the first to the other is greater than 1. The ratio is determined by taking the chance densities $f_\theta(x)$ and $f_\theta(x')$ for the two hypotheses involved. But $f_\theta(x)/f_\theta(x')$ could be greater than 1, while $f_\theta^*(y)/f_\theta^*(y')$ is less than or equal to 1.

To take care of this problem, Hacking introduces the notion of an experimental density. If y is a function of x and it is values of x which are determined by measurement, the experimental density of y according to h_θ should be $f_\theta(x) = f_\theta^*(y)dy/dx$.

Hacking's experimental densities raise more problems than they solve. Let y be a strictly monotonic and differentiable function of x with continuous derivative. Assuming that we

have some method for measuring x , do we not also have a method of measuring y ? On what basis does Hacking think priority may be given to one method of measurement over another? Hacking offers no explanation.

Hacking does extend his requirement that experimental densities be used to cases where a chance density distribution is defined for a variable which is a function of both the parameter θ and random variable x . Thus, if $v(x, \theta)$ satisfies the requirements for a pivotal function, Hacking insists that the likelihood ratio for d_v to $d_{v'}$ be given by $f_\theta(x)/f_{\theta'}(x)$, where $v = v(x, \theta)$ and $v' = v(x, \theta')$ rather than $g(v)/g(v')$. This is so regardless of whether the likelihood ratio is unconditional or is conditional on e_x . When the ratio is conditional, we can at least say relative to what value of x the experimental densities are given. But when e_x is not given, as when the likelihoods are taken to be unconditional, there is no uniqueness in the likelihood ratio except under special circumstances.

$g(v)\partial v/\partial x = f_\theta(x)$ for $v = v(x, \theta)$. If $\partial v/\partial x = k(v)w(x)$, then $g(v)/g(v') = f_\theta(x)/f_{\theta'}(x)$ no matter what value of x is picked. This condition is met by smoothly invertible and irrelevance-allowing pivotal functions.

But this entire exercise is utterly unnecessary. There is no need to consider likelihood distributions for values of v . We obtain Q -distributions for the hypotheses d_v by the principle of direct inference which Hacking endorses anyhow. There is no need to apply the law of likelihood.

As shall be explained shortly, Hacking invokes experimental densities in formulating his principle of irrelevance for the continuous case. Hence, we have another motive for considering them. But as I shall show shortly, the principle of fiducial inference for the continuous case may be derived from objectivist inductive logic and the law of likelihood for the continuous case as I have formulated it without appealing to the principle of irrelevance.

The formulation of the law of likelihood introduced here makes no use of experimental densities. Instead, I do not require that a likelihood ratio be uniquely determined for any pair of elements of W . I assume that such ratios are relative to a representation of the hypotheses in W by points in the joint space $\Omega \times \Theta$. If there is a strictly monotonic and differentiable transformation of the points in Ω with continuous

derivative, there must be a corresponding transformation of the *HL*-density function. With this proviso in place, a formulation of the law of likelihood no stronger than Hacking's (at least insofar as Hacking's formulation is clear at all) is obtainable. Moreover, it is demonstrable that this formulation suffices for the purpose of deriving the fiducial argument in the continuous case. This will now be proven.

**15.7
Likelihood and
Fiducial Infer-
ence in the
Continuous
Case**

Let the sample space be all points on some interval of the real line and let there be a smoothly invertible and irrelevance-allowing pivotal function satisfying condition (f^*) of chapter 14.

Because the pivotal function is irrelevance-allowing, the parameter space, sample space, and space of pivotal values can be so recast that $v(x, \theta) = v = x + \theta$ and each of three variables ranges from $-\infty$ to $+\infty$. There is no loss of generality, therefore, in considering the variables in a form satisfying this relation.

Consider $\eta(x, \theta; d_v)$ for some fixed value of v . We have $\eta(x, \theta; d_v) = f_\theta(x) = g(v)\partial v/\partial x$. But $\partial v/\partial x = 1$. Hence, $\eta(x, \theta; d_v) = g(v)$ as x varies over its entire range.

By the law of likelihood for the continuous case, for every $Q \in C(K)$, Q conditional on d_v is representable by a density function $f_v(x) = a$ constant k for all values of x . (This density will be a σ -finite density because the range of x is over the entire real line.) For all points (x, v) where v is held at the given fixed value, the joint density must be $f_v(x)g(v) = kg(v)$, which is a constant. The joint density for points (x, θ) such that $v(x, \theta) = v$ for the fixed value of v must be $f_\theta(x)h(\theta) = kg(v)\partial v/\partial x$, where $\partial v/\partial x = 1$. Hence, $h(\theta) = k$, which is a constant independent of the value of θ .

Consider now all elements in W consistent with K and e_x for some fixed value of x . Each of these is representable by a point (x, v) with density $kg(v)$. Hence, the conditional σ -finite density $g_x(v)$ must be directly proportional to $g(v)$. Since, however, $\int_{-\infty}^{\infty} g(v)dv = 1$, we can let $g_x(v) = g(v)$ and represent the conditional Q -distribution as a Q -density.

But this step is the adoption of confirmational irrelevance—the critical step of fiducial inference where a new principle of inductive logic is introduced. This step has been derived by adding the law of likelihood to our arsenal.

**15.8
The Principle
of Irrelevance
in the Contin-
uous Case**

Hacking proceeds in the continuous case as he does in discrete cases by invoking a principle of irrelevance.¹¹ Let V be a set of simple joint propositions each consistent with K and such that K entails that at least and at most one is true. Let V be indexed by values of a parameter taking all real values in some interval on the real line.

Let the experimental density values assigned to hypotheses in V equal the experimental density values assigned to hypotheses in V conditional on f (save for a positive constant factor). In that case, Hacking claims that the Q -distribution conditional on f should equal the unconditional Q -distribution for hypotheses in V except for a positive constant factor.

Thus, if the hypotheses in V are specifications of values of the pivotal $v = v(x, \theta)$, the experimental density for v is $f_\theta(x) = g(v)\partial v/\partial x$. This density function is not uniquely determined without specifying a value for x . If, however, the pivotal is irrelevance-allowing, then $\partial v/\partial x = k(v)w(x)$ so that $f_\theta(x)$ is a positive and constant multiple of $f_\theta(x')$. The likelihood of h_θ conditional on e_x is given by $f_\theta(x)$ for the specific value of x . The conditional likelihood function is a positive constant multiple of the unconditional likelihood function for v . Hence, the principle of irrelevance mandates the step of confirmational irrelevance when the pivotal function is smoothly invertible and irrelevance-allowing.

I have already explained my reservations concerning experimental densities. We now see why they appear so important in Hacking's argument. But as has been shown, they are not needed; for his principle of irrelevance can be dispensed with. The law of likelihood suffices as an addendum to objectivist inductive logic for the purpose of rationalizing fiducial reasoning.

**15.9
The Law of
Likelihood and
Indifference**

In the previous chapter, it was shown that the principles of fiducial inference, when added to objectivist inductive logic, must imply commitment to numerically definite prior credal distributions over the parameter space Θ . This is an inevitable consequence of taming fiducial inference in a Bayesian manner.

Hacking's law of likelihood even in the weakened versions I have used suffices to derive the three principles of fiducial inference introduced in the previous chapter. Hence, in cases

covered by those principles, Hacking's new principle of inductive logic mandates numerically precise prior credal states.

Hacking rather grudgingly acknowledges the commitment to priors at least in discrete cases but he seems to deny the commitment in continuous cases and, in any case, shrugs off such commitments with various obscure remarks.

Thus, he points out that in fiducial inference there is no need to derive posterior distributions explicitly from priors via Bayes' theorem.¹²

Of course this is true. But the force of his remark is not clear. The crux of the matter is that the law of likelihood entails a commitment to numerically precise priors regardless of how calculations are explicitly made. Euclidean geometry is axiomatizable in many ways. Adopting one set of axioms does not absolve one from commitment to the others.

Hacking seems to think, however, that he has a more substantial point to make than this. He correctly notes that if a pivotal function is smoothly invertible and irrelevance-allowing, all variables are transformable so that $v = x + \theta$ and all three variables range over the entire real line. Moreover, the prior credal distribution for values of θ must be uniform. H. Jeffreys was willing to embrace this result even though assigning h_θ positive density implies that the total probability is infinite.

In section 5.11, note was taken of Hacking's correct objection to Jeffreys' procedure in assigning h_θ positive density. But Hacking thinks that the only coherent option to Jeffreys' practice is to refuse to define a prior credal distribution altogether.¹³

But, in this case, Hacking's refusal is equivalent to assigning a uniform σ -finite distribution representing a Q -distribution violating countable additivity.

Hacking does not recognize this alternative and equivalent way of representing his view. Had he done so, he would have been compelled to recognize an inconsistency in his system; for, as noted previously, Hacking embraces Kolmogoroff's axioms for probability including countable additivity.

Disavowing the implications of one's commitments does not make them disappear. In section 12.16, an illustration of how contradictions slip in when the question of countable additivity is not kept under critical control was given. The contradiction

is obtained by adopting a credal state allowing violations of countable additivity (without acknowledging this) and then using countable additivity in the derivation of conclusions. The parallel between the arguments of M. Stone and his associates and Hacking's approach is inescapable—especially the inconsistency.

Of course, it is easy to eliminate the inconsistency. Either give up the law of likelihood or countable additivity. I think countable additivity should not be mandated as I have already explained. If we are to take Hacking's theory seriously, I think it best to remove countable additivity from his system as well.

The upshot, however, is that Hacking has failed to establish that his approach avoids commitment to numerically precise priors.

It is one thing to insist that Hacking's theory commits him to methods for assigning numerically precise priors. It is quite another to hold that he is committed to some principle of indifference mandating uniform prior credal distributions in states of ignorance.

Indeed, in my opinion, Hacking offers an interesting alternative to Jeffreys' (and Jeffreys' followers') prescription of uniform distributions of certain kinds in states of ignorance.

As I understand it, Hacking's thesis is that the law of likelihood is to be endorsed as a principle of inductive logic because it captures some conceptually necessary connection between chance and support. Recall that in defending principles of direct inference as principles of inductive logic, I too contended that unless such principles were introduced, we would lack a characterization of the connection between knowledge of chances and judgments about test behavior needed to provide chance with a meaningful role in inquiry and deliberation.

Hacking appears to understand the law of likelihood in the same vein. It is in this spirit that I interpret the following comment by Hacking on the relation between his principle of irrelevance (which I replaced by the law of likelihood) and principles of indifference.

The principle of irrelevance has been stated in complete generality. It makes no mention of special conditions of any problem. Its central notion, that degrees of support are as relative likelihoods, originates in the simplest conceptions about fre-

quency. It seems merely to continue that explication of the fundamental fact about frequency and support which led to the laws of likelihood. And even then the principle of irrelevance does not produce numerical degrees of support *de novo*, but merely, as it were, computes new probabilities out of those already established by the conventional frequency principle. . . .

Thus if the principle of irrelevance is true at all, it seems to express a fundamental fact about frequency and support. But turning to alleged judgements of indifference, we find no fundamental fact to express.¹⁴

If we take this strand in Hacking's thought seriously, it seems quite unnecessary for him to deny or hedge on the fact that his inductive logic entails a commitment to numerically precise priors under certain conditions. Even objectivist inductive logic allows numerically precise priors under some conditions. To be sure, Hacking's theory mandates such priors even in cases where objectivist logic does not; but if his argument is successful, his case for going beyond objectivism is congenial with views which spawn objectivism.

The sole ground for strengthening coherentist inductive logic to objectivist inductive logic is that principles of direct inference establish certain incorrigible links between knowledge of chance and judgments about test behavior where the incorrigibility is based on the claim that without such links the concept of chance would have no useful role in inquiry and deliberation.

If someone who accepts this argument is prepared to defend further principles on the grounds that they too impose restrictions on the connections between chance and test behavior needed in order that the concept of chance may be used significantly in inquiry and deliberation, the dispute between him and objectivists becomes a disagreement over details within the framework of important shared agreements.

Both parties to the dispute would agree that principles of inductive logic which are not introduced to establish needed links between cognitive attitudes towards chance hypotheses and hypotheses about test behavior are unacceptable—with the important exception of the principle of credal coherence, which has been defended on other grounds.

There is no need for excessive dogmatism. Perhaps some compelling argument may be given for adducing additional

principles of inductive logic. I do not know what such arguments could be, except for the fact that one can try and support anything one likes by claiming that it is "intuitive." I think we should minimize reliance on intuition and restrict its appeal to what is fairly noncontroversial.

I have not argued that Hacking's law of likelihood is, indeed, acceptable as a new principle of inductive logic. I do, however, think that his approach to defending it is along the right lines. But right lines or no, is his defense successful?

Thanks to T. Seidenfeld, it is demonstrable that Hacking's argument must fail as must any rationalization of fiducial inference along Bayesian lines. To this we now turn.

15.10 Seidenfeld's Theorem

Seidenfeld has shown that the principle of fiducial inference in the continuous case leads to contradiction.¹⁵ The implications of this result are of some interest. In the first place, the inconsistency arises when conditions are satisfied which D. V. Lindley claimed to have shown to be necessary and sufficient for the consistency of fiducial inference in the one-parameter case.¹⁶ Secondly, it has been shown here that a weakened version of the law of likelihood when grafted on to objectivist inductive logic entails the principles of fiducial inference. Consequently, Seidenfeld's result shows that either objectivist inductive logic must be abandoned or the law of likelihood modified. There seems to be no compelling way to tame fiducial inference along Bayesian lines and within the spirit of the objectivist view.

Let us now turn to Seidenfeld's argument.

X seeks to determine the volume of a hollow cube. He has two methods for doing so.

Method 1: The cube is filled with liquid of known density. The liquid is removed and weighed on a balance whose readings are normally distributed with unit variance and mean equal to the weight of the object measured. x_1 is the reading for weight and θ_1 the true weight of the liquid. The chance distribution for x_1 according to h_{θ_1} is

$$(1) f_{\theta_1}(x_1) = \frac{1}{\sqrt{2\pi}} e^{-(x_1 - \theta_1)^2/2}.$$

Let d_1 be the density of the liquid; $d_1 x_1 = y_1$ is the measured volume of the liquid and, hence, of the cube as revealed by the reading x_1 . The true volume is $w = d_1 \theta_1$.

$v_1 = x_1 - \theta_1$ is a smoothly invertible pivotal which is irrelevance-allowing and satisfies (f^*). The principle of fiducial inference implies that

$$(2) \quad h_{x_1}(\theta_1) = f_{\theta_1}(x_1).$$

The posterior distribution for w on the datum x_1 or y_1 is

$$k_{y_1}(w) = h_{x_1}(\theta_1) \frac{d\theta_1}{dw} = \frac{1}{d_1} h_{x_1}(\theta_1).$$

Consequently,

$$(3) \quad k_{y_1}(w) = \frac{1}{d_1} f_{\theta_1}(x_1)$$

holds.

Method 2: A metal bar is laid off against one edge of the cube and a length matching the length of one side of the cube is cut off. The bar has a constant cross-sectional area and constant density so that the length cut off is equal to a constant times its weight. If the true weight of the bar is θ_2 and the constant is d_2 , the true volume of the cube is $w = (d_2 \theta_2)^3$.

The piece of bar cut off is weighed on the same balance as before and a reading x_2 is obtained. The chance distribution for x_2 according to h_{θ_2} is

$$(4) \quad f_{\theta_2}(x_2) = \frac{1}{\sqrt{2\pi}} e^{(x_2 - \theta_2)^2/2}.$$

Let y_2 be a reading for the volume of the cube obtained from x_2 , so that $y_2 = (d_2 x_2)^3$.

$v_2 = x_2 - \theta_2$ is a smoothly invertible irrelevance-allowing pivotal satisfying condition (f^*). The principle of fiducial inference implies

$$(5) \quad h_{x_2}(\theta_2) = f_{\theta_2}(x_2).$$

Since $d\theta_2/dw = 1/(3d_2^3 \theta_2^2)$, then

$$(6) \quad k_{y_2}(w) = \frac{1}{3d_2^3 \theta_2^2} f_{\theta_2}(x_2).$$

Suppose that X makes a determination of the volume via method 1 and obtains the fiducial distribution (3) for w . He then contemplates making a further measurement using method 2. The distribution for values of w relative to the information obtained from both measurements is $k_{y_1 y_2}(w)$. It should be obtainable via Bayes' theorem using as a prior distribution for w the fiducial distribution (3) and as likelihoods the values for $t_w(y_2)$ where this density satisfies the condition

$$(7) \quad t_w(y_2) = \frac{1}{3d_2^3 x_2^2} f_{\theta_2}(x_2).$$

Appealing to (7) and (3), we conclude that

$$(8) \quad k_{y_1 y_2}(w) \propto \frac{1}{3d_1 d_2^3 x_2^2} f_{\theta_1}(x_1) f_{\theta_2}(x_2).$$

Suppose X had proceeded in reverse order. First he applied method 2 and had obtained x_2 and, hence, y_2 and the fiducial distribution $k_{y_2}(w)$. Using this distribution as the prior, he then employed Bayes' theorem to derive a new posterior $k_{y_2 y_1}(w)$ which must be directly proportional to the fiducial distribution $k_{y_2}(w)$ times $t_w(y_1)$. But the latter is

$$(9) \quad t_w(y_1) = \frac{1}{d_1} f_{\theta_1}(x_1).$$

From (9) and (6), we can conclude that

$$(10) \quad k_{y_2 y_1}(w) \propto \frac{1}{3d_1 d_2^3 \theta_2^2} f_{\theta_1}(x_1) f_{\theta_2}(x_2).$$

Compare (8) and (10). The right-hand sides are not only not equal but are not constant multiples of one another as w , and hence θ_2 , varies. Hence, the left-hand sides of (8) and (10) are not equal. But credal coherence requires that they be equal. We have a contradiction.

The contradiction derives from the principle of fiducial inference when combined with objectivist inductive logic. It must, therefore, be implied by adding to objectivist inductive logic any principle implying the principle of fiducial inference. Hacking's law of likelihood as formulated here has this property. So does his principle of irrelevance. Not only does the reconstructed version of his theory developed here lead to contradiction; it seems that his own version cannot be recon-

structed in any way reasonably matching what he claimed in his text without leading to the same contradiction.

It is easy to see the source of the inconsistency within the framework of Bayesian theory. If the fiducial argument is used with method 1, the prior for θ_1 is uniform (albeit an "improper" prior violating countable additivity). Hence, so is the prior for w , which is equal to the prior (density) for θ_1 divided by the constant d_1 .

If the fiducial argument is used with method 2, the prior for θ_2 is uniform (albeit improper). If the σ -finite density is k , the prior density for w becomes $k/(3d_2^3\theta_2^2)$, which is not uniform.

Those who have advocated principles of insufficient reason have had to face the problem of identifying for any given parameter θ which transformation to use in introducing a uniform distribution. Should the distribution of the unknown variance of a normal distribution prior to experimentation be uniform over values of the variance, the standard deviation, log variance, the precision, or what? The same applies to the mean and to every other continuous parameter. Unless some method for selecting a transformation of the parameter to use in assigning a uniform prior is given, principles of insufficient reason or indifference threaten to become inconsistent in the continuous case.

This is not the only difficulty with such principles; but it is a difficulty.

D. V. Lindley seemed to have shown that in those cases where fiducial arguments in one-parameter continuous cases avoid difficulties, there are definite criteria for identifying the appropriate transformations.

Lindley's results are correct provided one focuses attention only on data resulting from repetitions of the same kind of experiment.

The importance of Seidenfeld's example is that the data used to "estimate" the value of the unknown parameter (here, the volume of the cube) come from diverse experiments. He showed that one cannot escape the inconsistencies plaguing indifferentists by insisting on fiducialist restrictions.

But it is not only fiducialism that is undermined. Hacking's law of likelihood purports to formulate a principle of inductive logic that renders likelihood even more important to the understanding of chance and inverse inference than objectivists

have been prepared to acknowledge. Because, as has been shown here, the law of likelihood, when combined with objectivist inductive logic, leads to fiducial inference, Hacking's project of making likelihood the centerpiece of the logic of support has been decisively refuted.

15.11 Likelihood and Fiducial Probability: Another View

Before turning to the Kyburg-Dempster alternative to a Bayesian taming of fiducial inference, mention should be made of a method of linking likelihood and fiducial probability different from Hacking's.

It has already been acknowledged that both from the point of view advocated here and the strict Bayesian outlook, likelihoods are useful as indices of the contribution of information obtained from experimentation and observation to the determination of posterior credal states. Given the likelihood function on the data and the prior credal state, the posterior credal state can be uniquely determined by Bayes' theorem.

There is one minor limitation on the use of likelihoods in this way. Likelihoods are not defined for composite chance hypotheses on the data.

It is possible to use fiducial distributions as surrogates for likelihoods in "summarizing" the contribution of data to both composite and simple chance hypotheses when fiducial distributions exist. On this interpretation, fiducial probabilities are not construed as posterior credal probabilities but as factors which along with prior probabilities determine the posterior credal state.

Let $v = x + \theta$, and construct the fiducial distribution $h_x(\theta) = g_x(v) = g(v)$. Let the posterior distribution $h_x^*(\theta)$ be determined from this by setting

$$h_x^*(\theta) = \frac{h_x(\theta)h^*(\theta)}{\int_{-\infty}^{+\infty} h_x(\theta)h^*(\theta)d\theta}$$

where the prior is $h^*(\theta)$. Needless to say, if this prior distribution is uniform, the posterior distribution equals the fiducial distribution.

What we have done, in effect, is to regard the posterior as directly proportional to the fiducial density times the prior density.¹⁷

Thus, we have a construal of fiducial probability which both Bayesians and those who are sympathetic to the proposals

made in this book can accept. (Notice that one can use the fiducial distribution with indeterminate priors to obtain indeterminate posteriors.) Moreover, fiducial probability and likelihood are understood as serving the same function. These aspects of the interpretation are congenial to the views of both Fisher and Hacking.

However, both of these authors take fiducial probabilities to be probabilities on data in the same sense in which posteriors computed via Bayes' theorem are. This interpretation of fiducial probability cannot, therefore, be the one they have been looking for.

16.1 Fisher on Direct Inference

R. A. Fisher always denied that fiducial inference entails a commitment to numerically precise priors. Moreover, in his later writings, he took the position that the principles of direct inference and confirmational coherence did not have to be supplemented with additional principles in order to justify fiducial inferences. Fisher may not have adhered to these claims consistently; but he did seem to endorse them. To this extent, the efforts to tame fiducial inference of the sort described in chapters 14 and 15 cannot pretend to be a reconstruction of Fisher's view.

There are two approaches to deriving probability judgments from data which appear to come closer to Fisher's view. H. E. Kyburg and A. P. Dempster have offered accounts of how to assign numerically precise probabilities to statistical hypotheses on the basis of data which do not entail a commitment to numerically precise priors. Kyburg's approach is explicitly based on a commitment to confirmational coherence and a principle of direct inference. Dempster's approach permits a similar interpretation.

Such views must conflict somehow with the conditions on confirmational commitments entailed by an objectivist outlook. Kyburg and Dempster recommend rejecting confirmational conditionalization. Kyburg does so by modifying the principle of direct inference in a manner which entails such rejection. As I reconstruct his view, Dempster allows the principle of direct inference to stand more or less as I have formulated it but recommends a revision of the principle of confirmational conditionalization.

Fisher's statements are far too enigmatic to warrant a decisive verdict as to who captures his intent best. I suspect that Kyburg's view comes closest to Fisher's.

Perhaps the clearest statement of Fisher's view is found in "Mathematical Probability in the Natural Sciences,"¹ al-

though he develops the same themes in *Statistical Methods and Scientific Inference*² and other writings of the late 1950s.³ In that essay, he claims that there are three "requirements of a correct statement of Mathematical Probability."⁴

The first of these he calls a "conceptual Reference Set which may be *conceived* as a population of possibilities of which an exactly known fraction possess some chosen characteristic."

Thus, a "correct statement of Mathematical Probability" is based on knowledge. Such a statement may, therefore, be construed as an evaluation of a hypothesis with respect to credal probability relative to appropriate knowledge.

Moreover, the knowledge specified by the first condition is of the percentage of members of a "conceptual Reference Set" which "possess some chosen characteristic." It is quite clear from Fisher's own elaborations that he does not intend literally that X know the percentage of events or objects having some characteristic in a conceptual reference class. Nor does he mean that the knowledge be of the relative frequency in some population of past, present, and future objects or events.

However, the probability statements do not imply the existence of any such population in the real world. All that they assert is that the exact nature and degree of our uncertainty is just *as if* we knew it to have been chosen at random from such a population.⁵

Unfortunately, Fisher fails to explain the import of the "as if" in this passage. It does seem clear, however, that the knowledge X has (not the "as if" knowledge) is analogous to knowledge (which X may sometimes have) that a certain population contains members with a given trait with a definite relative frequency and that the chance of obtaining an object having that trait on a selection from that population of a certain kind called a "random selection" is equal to the percentage of objects in that population having that trait.

This suggests that the knowledge X has is knowledge that a given chance setup has a simple chance property. Fisher does not interpret (in the sense of giving satisfaction conditions for) the chance predicate in terms of relative frequencies in a conceptual reference class but, as R. B. Braithwaite observed,⁶ uses the conceptual reference class as a model for heuristic purposes. Fisher is most emphatically not a partisan

of "frequency interpretations" of either chance (statistical probability) or of the probability of hypotheses (credal or confirmational probability). This is clear both from the passage cited and from his chastisement of J. Neyman in the same essay.⁷

On this analysis, Fisher's first necessary condition for a "correct statement of Mathematical Probability" is knowledge that the chance of obtaining a result of kind R on a trial of kind S on chance setup a is some definite value r .

Fisher's second requirement is that "it must be possible to assert that the subject of the probability statement belongs to this set"—i.e., the conceptual reference set.⁸ Within the framework I am using, this condition amounts to specifying that X know that the specific trial on chance setup a be of kind S .

The third requirement, which Fisher insists is the most interesting, is that "no subset can be recognized having a different probability. Such subsets must always exist; it is required that no one of them shall be recognizable. This is a postulate of ignorance, and therefore unfamiliar to deductive reasoning, though a characteristic of inductive logic."⁹

The third requirement is intended to serve a function similar to the proviso in my principle of direct inference which states that extra information known to be true of the trial of kind S should be known to be stochastically irrelevant. However, it is not clear that the third condition is equivalent to this proviso.

To explain the obscurity, I assume that Fisher's three "requirements" are not only individually necessary but jointly sufficient for a "correct statement of Mathematical Probability"—i.e., for an assignment of a numerically precise degree of credal probability to a hypothesis about the outcome of a trial on a chance setup. In other words, the three requirements are intended as a principle of direct inference combined with the stipulation that numerically precise credal states are legitimate only when they are derivable via direct inference from knowledge of chances.

The three conditions become, therefore, requirements on X 's corpus of knowledge K individually necessary and jointly sufficient for assigning a numerically precise degree of cre-

dence to a hypothesis about the outcome of a trial relative to that corpus.

Let us rehearse the three conditions:

- (i) *K* contains "The chance of an *R* on a trial of kind *S* is equal to *p*."
- (ii) *K* contains a statement asserting that at some time *t* a trial of kind *S* occurs.
- (iii) *K* does not contain a statement asserting that the trial of kind *S* at *t* is also of kind *T* where *K* also contains "The chance of an *R* on a trial of kind *S* & *T* differs from *p*."

Fisher proceeds to illustrate condition (iii).¹⁰ He cites situations where *K* contains information that the trial of kind *S* is of kind *T* as well but also contains the assumption that the chance of an *R* on a trial of kind *S* & *T* is equal to *p*. In that event, *X* knows that the extra information is stochastically irrelevant and, hence, may be ignored.

Fisher also considers cases where the trial is of kind *T* as well as of kind *S* but is not known to be of kind *T*. In agreement with my version of direct inference, he recommends that *p* be assigned as the degree of credence of the hypothesis that an event of kind *R* occurs.

What about situations where *K* contains the information that the trial is of kind *T* as well as of kind *S* but lacks information concerning the precise chance of an *R* on a trial of kind *S* & *T*? As far as I can discover, Fisher nowhere explicitly deals with cases of this sort.

Suppose, however, that the result in question were a pivotal event and the trial of kind *S* & *T* were a trial of kind *S* yielding a certain value for the random variable *x*. This is an example where the agent has knowledge of chances on trials of kind *S* (concerning pivotal values) but lacks such knowledge on trials of kind *S* & *T*. Fisher seems to think that direct inference alone mandates a degree of credence equal to the chance on a trial of kind *S*. That is to say, Fisher seems to hold that at least sometimes such extra information *T* may be ignored.

If this is Fisher's intent, his principle of direct inference must differ from mine. According to the principle I have proposed, unless *X* knows the extra information to be stochastically irrelevant, he should base his direct inference on his

knowledge of the chance of an *R* on a trial which is of kind *S* & *T*. If he does not know the chance precisely but only knows a composite chance assumption, the degree of credence assigned will depend on his credal probability state over the rival chance hypotheses ingredient in the composite chance hypothesis as explained in sections 12.13-12.15.

For some such credal states, the degree of credence assigned the hypothesis that an *R* occurs will equal *p*. But this presupposes commitment to the confirmational irrelevance of the extra information. That assumption of confirmational irrelevance is not derivable from the principle of direct inference alone.

Fisher's practice and some of his remarks suggest that on some occasions when the extra information about the kind of trial is not known to be stochastically irrelevant, the principle of direct inference should imply confirmational irrelevance. Unfortunately, Fisher fails to give us any explicit formulation of the circumstances when this is so.

H. Reichenbach proposed an account of direct inference according to which one should base the direct inference on the narrowest reference class for which one had "reliable statistics."¹¹ Reichenbach failed to explain what he meant by reliable statistics. I think he meant that one should have evidence warranting a precise judgment of chance or relative frequency in the reference class.

Thus, if *X* knows that the chance of an *R* on a trial of kind *S* is *p* but also knows that the trial is of kind *S* & *T*, *X* should assign a degree of credence to the hypothesis that an *R* will occur equal to *p* if he knows nothing as to the chance of an *R* on a trial of kind *T* or on a trial of kind *S* & *T*. In that event, the reference class of trials of kind *S* is the "narrowest" for which *X* has "reliable statistics" (precise information?).

On the other hand, if *X* has reliable statistics for *S* and for *T* but not for *S* & *T* and the chance of an *R* on a trial of kind *S* differs from that for a trial of kind *T*, there is no unique narrowest reference class for which *X* has reliable statistics. Reichenbach acknowledges his prescriptions are ambiguous in such a case. The best he can do is suggest collecting further statistics about trials of kind *S* & *T*.

Of course, this advice cannot always be followed. Reichenbach does not say anything about "weight" or credence when

one cannot collect the extra statistics. I think it fair to say that in situations of the sort just envisaged, the degree of credence assigned to the hypothesis that an R occurred is the maximum interval 0 to 1.

Neither Reichenbach nor Fisher deal with these cases in anything like a systematic or exhaustive fashion. To my knowledge, the only one to do so is Kyburg. This circumstance is, in my opinion, scandalous, considering the considerable attention which has been devoted by philosophers to the reference class.

To be sure, one can proceed as I have done and require that extra information must be known to be stochastically irrelevant if it is to be ignored on the grounds of confirmational irrelevance. Such an approach is not, so I believe, Fisher's or Reichenbach's. One of the reasons (by no means the only one) why Kyburg's ideas are important to consider is that he alone has offered a systematic account of direct inference where confirmational irrelevance of extra information about kinds of trials is mandated by direct inference even though knowledge of stochastic irrelevance is absent.†

As shall be explained, if this rival version of the principle of direct inference is adopted, it is, indeed, possible to secure the legitimacy of fiducial inferences without principles of inductive logic additional to confirmational coherence and the Kyburgian version of direct inference. Furthermore, commitment to numerically precise prior credal states is avoided.

† Kyburg's principle of direct inference is embedded in his definition of an object being a "random member" of a set and in his definition of epistemological probability. The theory is presented in *Probability and the Logic of Rational Belief* (Middletown, Conn.: Wesleyan University Press, 1961), especially ch. 9 and *The Logical Foundations of Statistical Inference* (Dordrecht: Reidel, 1974), especially chs. 9-11. Because I offer no definitions of credal or confirmational probability but, instead, impose conditions on confirmational commitments and credal states and wish to explain Kyburg's view within this framework, my presentation of his view deviates from his in formulation. The notion of K -irrelevance corresponds roughly to his notion of random membership and his conception of epistemological probability relative to a corpus corresponds to the weakest or logical credal state relative to the corpus when the principle of direct inference I attribute to Kyburg is substituted for the one I advocate. Otherwise, the chief difference is that he takes direct inference to be based on knowledge of frequencies whereas I take it to be based on knowledge of chances. This latter difference is taken up in section 16.6.

Thus, not only is the Fisherian idea of direct inference captured in a precise form; but Fisher's hopes for fiducial inference seem open to realization. The difficulties with objectivist necessitarianism are surmounted by adopting a principle of direct inference alternative to the one I have proposed.

16.2 Kyburg's Theory

Kyburg's account of direct inference is far more general than anything found in Fisher or Reichenbach. This is due not only to the manner in which he handles the so-called problem of the reference class, but also due to the way he treats direct inference from knowledge of composite chance assumptions. Both Fisher and Reichenbach restrict their accounts to situations where direct inference is based on knowledge of simple chance assumptions.

The principle of direct inference I have proposed is also formulated to apply to cases where the knowledge is of simple chance hypotheses. On my account, direct inference from knowledge of composite chance assumptions relies on confirmational coherence and conditionalization as well as on direct inference. Kyburg's principle of direct inference violates confirmational conditionalization in a manner preventing him from handling inferences from knowledge of composite chance assumptions as I have done. Consequently, he formulates his principle in a manner applicable both to cases where the inference is based on knowledge of composite chance assumptions and where the knowledge is of simple chance assumptions.

Kyburg is one of the first authors, after B. O. Koopman, I. J. Good, and C. A. B. Smith, to develop an account of interval-valued credal or confirmational probability and, in many respects, he pushed his approach more deeply and systematically than these authors did.

According to Kyburg, probability is an "epistemological" notion. This means that probability, as he understands it, is relative to a corpus of knowledge. Such a notion corresponds to my notion of a confirmational commitment legislated by principles of inductive logic alone.

In this sense, Kyburg's notion of probability resembles that considered by J. M. Keynes and H. Jeffreys and the one on which R. Carnap focused most of his attention. In effect, his epistemological probabilities are those credal probability judg-

ments determined by the logical confirmational commitment—i.e., the weakest allowed by the principles of inductive logic.

Kyburg allows the credal state to assign numerically indeterminate but interval-valued degrees of credence to hypotheses on the evidence. His principle of direct inference not only covers cases where the chance assumptions are composite but where the credal states generated by knowledge of chances are numerically indeterminate.

Another salient feature of Kyburg's view is his rejection of the concept of statistical probability or chance. On his view, direct inference is based on knowledge of frequencies. Unlike Fisher, he does not appeal to "conceptual" or "hypothetical" reference classes in a metaphorical manner. His reference classes are intended to be sets of events or objects. Instead of using knowledge of chances, he appeals to knowledge of relative frequencies with which traits appear or occur in such populations.

For the present I shall ignore this feature of Kyburg's approach, returning to it in section 16.6. No distortion of Kyburg's account of direct inference results. That account applies to direct inferences based on knowledge of chances or inferences based on knowledge of relative frequencies in populations.

Let X know the following:

- (i) The chance of an R on a trial of kind S falls in the interval from \underline{p} to \bar{p} .
- (ii) The chance of an R on a trial of kind T falls in the interval from \underline{p}' to \bar{p}' .
- (iii) The chance of an R on a trial of kind S & T falls in an interval from \underline{q} to \bar{q} .

In all cases, the intervals specified are the narrowest in which the chance is known to fall.

X also knows that the trial at some time is both of kind S and of kind T .

To illustrate, let a trial of kind Z be the selection at random from the population of persons who are either Swedes or visitors to Lourdes and let the result of kind R be obtaining a Protestant on such a selection.

The chance of obtaining a Protestant on such a selection is then equal to the percentage of Protestants in that population.

A trial of kind S is a trial of kind Z which yields a Swede. The chance of an R on a trial of kind S is equal to the percentage of Protestants among the Swedes.

A trial of kind T is a trial of kind Z yielding a visitor to Lourdes. The chance of an R on a trial of kind T is equal to the percentage of Protestants among the visitors to Lourdes.

A trial of kind Z which is both of kind S and of kind T is one which yields a Swedish visitor to Lourdes. The chance of an R on a trial of that kind is equal to the percentage of Protestants among the Swedish visitors to Lourdes.

Relative to X 's corpus K as specified above, the information that a trial is of kind T is *L-irrelevant* to the issue as to whether an R occurs on that trial relative to information that the trial is of kind S if and only if K contains the information that being of kind T is stochastically irrelevant to yielding an R on a trial of kind S .

In our example, the information that the random selection yields a visitor to Lourdes is *L-irrelevant* to whether the person selected is a Protestant relative to information that the person is a Swede if it is known that the percentage of Swedish visitors to Lourdes who are Protestant is identical with the percentage of Swedes who are Protestants.

L-irrelevance is an epistemological notion specifying a feature of the corpus K . Stochastic irrelevance is not epistemological. The principle of direct inference I have proposed could be formulated as stating that all extra information about a trial of kind S be *L-irrelevant* rather than that it be known to be stochastically irrelevant. The concept of *L-irrelevance* is introduced here for the purely temporary purpose of comparing the principle of direct inference I have proposed with Kyburg's.

Relative to K , the information that the trial under consideration is of kind T is *K-irrelevant* to the issue as to whether an R occurs on that trial relative to information that the trial is of kind S if and only if X 's corpus K implies the following:

- (a) The interval $[\underline{p}, \bar{p}]$ falls within or coincides with the interval $[\underline{q}, \bar{q}]$.

(b) The interval $[\underline{p}, \bar{p}]$ falls within or coincides with the interval $[\underline{p}', \bar{p}']$.

S is the total information K -relevant to the occurrence of an R if and only if S is known to be true of the trial event and the following two conditions are satisfied:

(c) if S does not entail T (while T is known to be true of the trial event), T is K -irrelevant to the hypothesis that an R occurs with respect to S .

(d) if S entails S^* while S^* does not entail S , there is some T^* not entailed by S^* but known true of the trial event which is K -relevant to the hypothesis that an R occurs with respect to S^* .

Kyburg's principle of direct inference stipulates that if S is the total information K -relevant to the occurrence of an R , the degree of credence which should be assigned to the hypothesis that an R occurs relative to the corpus K should be $[\underline{p}, \bar{p}]$.

This principle differs from the principle of direct inference proposed in this essay in two respects: (i) the condition of L -irrelevance is replaced by the condition of K -irrelevance and (ii) when X 's knowledge of the chance of an R on a trial of kind S is composite, the degree of credence to be assigned the hypothesis that an R occurs is mandated to be the entire interval from \underline{p} to \bar{p} .

The difference of primary interest here is the first one. The second difference disappears in the case of objectivist necessitarianism.

To appreciate the significance of the first difference and how it bears on the views of Fisher and Reichenbach, I shall consider a series of different cases and compare Kyburg's approach to mine in each one.

For the sake of simplicity, I shall suppose the strongest information X knows about the trial in question is that it is of kind S & T .

Case 1: X knows

- (1) The percentage of Protestants among Swedes is .9.
- (2) The percentage of Protestants among visitors to Lourdes is .9.

(3) The percentage of Protestants among Swedish visitors to Lourdes is .9.

Being a trial of kind T is L -irrelevant and K -irrelevant with respect to being a trial of kind S ; and being a trial of kind S is L -irrelevant and K -irrelevant with respect to being a trial of kind T . Both Kyburg's and my principles of direct inference mandate assigning a degree of credence of .9 to the hypothesis that the person selected is a Protestant.

Case 2: X knows (1), (3), and

(4) The percentage of Protestants among the visitors to Lourdes is .1.

Being a trial of kind T is L -irrelevant but not K -irrelevant with respect to being of kind S . Being a trial of kind S is both L -relevant and K -relevant to being of kind T . In spite of the differences, both principles of direct inference mandate assigning a degree of credence of .9 to the hypothesis that the person selected is a Protestant.

Case 3: X knows (1), (3) and

(5) The percentage of Protestants among visitors to Lourdes is .85 or .95.

Being of kind T is both L -irrelevant and K -irrelevant with respect to being of kind S . Being of kind S is both L -relevant and K -relevant with respect to being of kind T . Both principles of direct inference mandate assigning a degree of credence of .9 to the hypothesis that the person selected is a Protestant.

Case 3a: X knows (1), (3), and

(5a) The percentage of Protestants among visitors to Lourdes is .85, .9, or .95.

Case 3b: X knows (1), (3), and

(5b) The percentage of Protestants among visitors to Lourdes is some value between 0 and 1.

In case 3, X knows that the chance of obtaining a Protestant on a trial of kind T is different from the chance of obtaining a Protestant on a trial of kind S . In cases 3a and 3b, X does not know this or its negation. The appraisals of K -irrelevance

and *L*-irrelevance remain the same as in case 3 and so does the verdict of direct inference.

Case 4: *X* knows (1) and one of (2), (4), (5), (5a), or (5b). He also knows

(6) The percentage of Protestants among Swedish visitors to Lourdes is .85.

Being of kind *T* is both *L*-relevant and *K*-relevant with respect to being of kind *S* and vice versa. The degree of credence assigned to the hypothesis that the individual drawn is Protestant is .85 according to both principles of direct inference.

Case 5: This is like any of the cases covered by case 4 except that in lieu of (6) *X* knows

(7) The percentage of Protestants among Swedish visitors to Lourdes is either .75 or .85.

Being of kind *T* is both *L*-relevant and *K*-relevant with respect to being of kind *S* and vice versa. According to Kyburg, the degree of credence assigned to the hypothesis that the individual drawn is Protestant should be the interval from .75 to .85. According to the principle of direct inference favored here, it should be some subinterval of that one. If objectivist necessitarianism is endorsed, the interval should be, as for Kyburg, from .75 to .85.

Although some small differences emerge in case 2 concerning evaluations of *K*- and *L*-irrelevance, in neither this case or the other four considered is there any difference concerning the results of direct inference.

Consider now the following hypotheses about the percentage of Swedish visitors to Lourdes who are Protestants:

(8) The percentage of Protestants among Swedish visitors to Lourdes is 85% or 95%.

(8a) The percentage of Protestants among Swedish visitors to Lourdes is 85%, 90%, or 95%.

(8b) The percentage of Protestants among Swedish visitors to Lourdes is anywhere from 0% to 100%.

Reichenbach considered a case exemplified by

Case 6: *X* knows (1), (4), and (8b).

In this case, *X* knows that the chance of an *R* on a trial of kind *S* is different from the chance of an *R* on a trial of kind *T* and knows nothing of the chance of an *R* on a trial of kind *S* & *T*. Reichenbach's view seems to be that the "weight" is undetermined in this case pending information about the chance of an *R* on a trial of kind *S* & *T*. That is to say, the degree of credence to be assigned the hypothesis that an *R* has occurred (a Protestant selected) is the interval from 0 to 1.

Both my analysis and Kyburg's give this prescription as well. Being of kind *T* is both *L*-relevant and *K*-relevant with respect to being a trial of kind *S* and vice versa. Kyburg's rule mandates the interval from 0 to 1. So does my rule when combined with necessitarianism. When necessitarianism is not mandated, some subinterval or other is required.

Reichenbach fails to consider variants of this case where *X* knows (1), (4), and some more precise information than conveyed by (8b) concerning the percentage of Swedish visitors to Lourdes who are Protestants. Case 4 provides an example, as do the variants on case 6 substituting (8) and (8a) for (8b). Let case 6a be the result of substituting (8) for (8b) and case 6b be the result of substituting (8a). Both my rule and Kyburg's prescribe assigning the interval-valued credence from .85 to .95. I think this is congenial with Reichenbach's outlook.

Reichenbach *seems* to suggest, however, that when the chance of an *R* on a trial of kind *S* is the same as the chance of an *R* on a trial of kind *T* and the chance of an *R* on a trial of kind *S* & *T* is not known—i.e., could range from 0 to 1—the degree of credence to be assigned to the hypothesis that a result of kind *R* occurs should equal the chance of an *R* on a trial of kind *S*. Whatever Reichenbach may have intended, however, Kyburg does explicitly prescribe this result.

Case 7: *X* knows (1), (2), and (8).

Case 7a: *X* knows (1), (2), and (8a).

Case 7b: *X* knows (1), (2), and (8b).

In all three cases, being of kind *T* is *L*-relevant with respect to being of kind *S* and vice versa. In all three cases, being of

kind T is K -irrelevant with respect to being of kind S and vice versa.

The principle of direct inference prescribes the interval from .85 to .95 in cases 7 and 7a. Kyburg's principle prescribes .9. In case 7b, the principle of direct inference prescribes the interval from 0 to 1. Kyburg's principle, once more, prescribes .9. My principle, of course, makes these prescriptions under the assumption of necessitarianism.

The differences between Kyburg's view (which, so I believe, captures Reichenbach's intent here) and my approach just exemplified is manifested in the following variations as well.

Case 8: X knows (1) and either (5), (5a), or (5b) and either (8), (8a), or (8b).

Being of kind T is L -relevant but K -irrelevant with respect to being of kind S . Being of kind S is both L -relevant and K -relevant with respect to being of kind T .

Kyburg's approach requires basing credence on knowledge of the chance of an R on a trial of kind S and, hence, recommends .9.

My approach requires basing credence on knowledge of the chance of an R on a trial of kind S & T and, hence, recommends the interval from .85 to .95 if (8) or (8a) are assumed and from 0 to 1 when (8b) is accepted.

Fisher was even less explicit than Reichenbach on the relevant cases. However, if we apply Kyburg's principle of direct inference to cases where fiducial inference is entertainable, Fisher's prescriptions are often mandated without adding any further principles of inductive logic. This will now be shown.

16.3 Kyburg's Theory and Fiducial Inference

Let $v = v(x, \theta)$ be a smoothly invertible and irrelevance-allowing pivotal satisfying (f) or (f^*) of chapter 14.

The prior credal state is assumed maximally indeterminate for the values of θ .

The credal distribution for v is given by $g(v)$ prior to finding out the result of experimentation. This is so according to Kyburg's theory and my own if, as I shall suppose, the trial is of kind S & T and being of kind T is L -irrelevant to being of kind S and, hence, is K -irrelevant as well.

An observation is made and the value x is discovered. But X does not know the chance distribution for the values v on a trial of kind S & T which yields that value x . That is to say, X does not know whether yielding value x is stochastically irrelevant or not. Hence, the information that the trial is one yielding the reading x is L -relevant information. According to my proposal, it cannot be ignored. An objectivist necessitarian is obligated to assign a credal distribution for values v conditional on e_x which is maximally indeterminate. To obtain a fiducial distribution, the principle of fiducial inference has to be invoked with its untoward results.

Observe, however, that although X does not know that the information that the trial is one yielding the reading x is stochastically irrelevant, he also does not know that it is stochastically relevant. Hence, it is K -irrelevant information.

Kyburg's theory, therefore, allows X —indeed, obligates X —to ignore the extra information and adopt $g_x(v) = g(v)$ as his conditional distribution for v . From this, a numerically precise distribution for θ emerges conditional on e_x as in fiducial inference.

The situation here is quite analogous to that in case 8.

Thus, Kyburg's theory succeeds in doing what Fisher claimed his account of fiducial inference does—to wit, it obtains fiducial posteriors without a commitment to numerically precise priors and with the aid of direct inference where the choice of a "recognizable subset" is the critical factor.†

However, Kyburg's success depends on his violating confirmational conditionalization. Moreover, that violation must be due to the version of direct inference he endorses. This point should be brought out more explicitly.

16.4 Kyburgian Direct Inference and Confirmational Conditionalization

Let e be the hypothesis that the individual selected is a Protestant. h_θ asserts that the percentage of Protestants among Swedish visitors to Lourdes is 100 θ %. X knows that 90% of Swedes are Protestants. He is utterly ignorant concerning the percentage of visitors to Lourdes who are Protestants; but he assumes that either $h_{.89}$, $h_{.91}$, or $h_{.92}$ is true.

† Kyburg's theory is able to allow posterior distributions which are only partially indeterminate beginning with maximally indeterminate priors. Fisher does not recognize such cases.

According to Kyburg's principle of direct inference, the following holds:

$$(1) Q(e) = .9.$$

If confirmational conditionalization is operative, Kyburg's principle implies the following conditions on Q -distributions in $C(K)$ as well:

$$(2) Q(e; h_{.89}) = .89$$

$$(3) Q(e; h_{.91}) = .91$$

$$(4) Q(e; h_{.92}) = .92$$

$$(5) Q(e; h_{.89} \vee h_{.91}) = .9$$

$$(6) Q(e; h_{.89} \vee h_{.92}) = .9$$

$$(7) .91 \leq Q(e; h_{.91} \vee h_{.92}) \leq .92.$$

Credal coherence implies that

$$(8) Q(e) = Q(e; h_{.89} \vee h_{.91})Q(h_{.89} \vee h_{.91}; h_{.89} \vee h_{.91} \vee h_{.92}) \\ + Q(e; h_{.92})Q(h_{.92}; h_{.89} \vee h_{.91} \vee h_{.92}).$$

From (1), (4), (5), and (8), we have

$$(9) .9 = .9Q(h_{.89} \vee h_{.91}; h_{.89} \vee h_{.91} \vee h_{.92}) \\ + .92Q(h_{.92}; h_{.89} \vee h_{.91} \vee h_{.92}).$$

Consequently,

$$(10) Q(h_{.92}) = Q(h_{.92}; h_{.89} \vee h_{.91} \vee h_{.92}) \\ = 0.$$

By similar reasoning,

$$(11) Q(h_{.91}) = 0.$$

Hence,

$$(12) Q(h_{.89}) = 1.$$

But then it must be the case that

$$(13) Q(e) = Q(e; h_{.89})Q(h_{.89}) \\ = .89.$$

The statement (13), however, contradicts (1).

Kyburg's earliest writings on the subject suggest that he might have been committed to confirmational coherence and

consistency.† On this reading of his view, he was committed to abandoning confirmational conditionalization.

In any case, he was committed to denying that all shifts in credal state where the confirmational commitment is held constant should guarantee that posterior credal states be derivable from the prior states via Bayes' theorem. More recently Kyburg has recognized this explicitly and has cited some examples of Bayesian breakdown.¹² He does not appear to have realized, however, how intimately this breakdown is linked with his version of the principle of direct inference.

The account of his view presented here renders this connection explicit and indicates how closely linked the problem of the reference class is to the question of the tenability of confirmational conditionalization.

Once confirmational conditionalization is abandoned in the manner just indicated, Seidenfeld's example of the hollow cube does not yield a contradiction according to fiducialist principles. The contradiction is obtained with the aid of the assumption of confirmational conditionalization. Kyburg's theory may be used to provide an analysis of this situation; but it is too complex to consider here.

16.5 Dempsterian Conditionalization

Kyburg's theory entails a violation of confirmational conditionalization due to his modification of the principle of direct inference. An account of confirmational commitments is envisageable that retains intact confirmational consistency, coherence, convexity, and direct inference, but that revises confirmational conditionalization.

A. P. Dempster has proposed an account of interval-valued probability suggesting such an outlook. I say "suggesting" because I am not quite sure what Dempster's view actually is, even though it is possible to develop an account of the sort

† See *Probability and the Logic of Rational Belief* (Middletown, Conn.: Wesleyan University Press, 1961), pp. 222-224. On these pages, Kyburg correctly claims that given an interval-valued credal state meeting his requirements there is at least one probability distribution assigning probabilities within these interval specifications satisfying coherence requirements. Confirmational coherence and consistency are, therefore, satisfied. On these pages and those preceding, he denies that his intervalist theory satisfies coherence requirements—but this means only that upper and lower probabilities are not probability measures and that, from Kyburg's point of view, credal states need not be numerically determinate.

just indicated from ingredients found in his writings.† The account to be presented is just such a reconstruction.

K_e is the expansion of K obtained by adding e and forming the deductive closure. Let the range of permissible values for e in $C(K)$ range from \underline{s} to \bar{s} .

$C(K_e)$ is the D -conditionalization of $C(K)$ with respect to K and K_e if and only if (i) B^* is the subset of $C(K)$ consisting of Q -distributions assigning e the value \bar{s} and (ii) $C(K_e)$ is the conditionalization of B^* with respect to K and K_e .

D -Confirmational Conditionalization: For every K and K_e where e is consistent with K , $C(K_e)$ is the D -conditionalization of $C(K)$ with respect to K and K_e .

If every permissible Q -value in $C(K)$ assigns e the same real value, $C(K_e)$ is the D -conditionalization of $C(K)$ if and only if it is the conditionalization of $C(K)$. Thus, the principle of D -confirmational conditionalization implies that the principle of confirmational conditionalization is applicable in special cases.

But D -conditionalization can be consistently applied only if further restrictions are imposed on the credal states allowed relative to potential corpus K . To explain this, consider the three propositions $e \ \& \ f$, $e \ \& \ \sim f$, and $\sim e$ and let the credal state $C(K)$ be the convex hull of the four distributions

- (i) 1/2, 1/8, 3/8 (ii) 3/8, 3/8, 1/4
 (iii) 1/8, 1/8, 3/4 (iv) 1/8, 1/2, 3/8.

Suppose the partition is refined by construing each of the three alternatives as a disjunction of hypotheses in which g is

† The version of Dempster's view I am going to discuss is based on "Upper and Lower Probabilities Induced by a Multivalued Mapping," *Annals Math. Stat.*, v. 38 (1967), pp. 325-339; and "A Generalization of Bayesian Inference," *J. Royal Stat. Soc.*, ser. B, v. 30 (1968), pp. 205-247 (with discussion). See also the somewhat different emphasis given in "New Methods for Reasoning Towards Posterior Distributions Based on Sample Data," *Annals Math. Stat.*, v. 35 (1966), pp. 355-374. My account of Dempster's theory relies more heavily than his does on representations of probability states by convex sets of distributions. Dempster himself makes use of this mode of representation (e.g., in "Upper and Lower Probabilities," pp. 330-331 and in "A Generalization of Bayesian Inference," p. 245). To my knowledge, Dempster nowhere formulates what I call D -confirmational conditionalization in the form I give; but he uses a principle of "conditioning" which appears to be equivalent to it. I think Dempster is committed to a version of direct inference of the kind I favor (as opposed to Kyburg's); but I am not confident of this.

true and in which g is false and consider the conditional distributions for g and $\sim g$ relative to the three hypotheses.

K_1 is obtained from K by adding e , K_2 from K_1 by adding f , and K^* from K by adding $e \ \& \ f$. Hence $K^* = K_2$. The difference lies only in the sequence of expansions involved in shifting from K to K^* .

In shifting from K to K_1 , D -conditionalization mandates going to $C(K_1)$ which is the ordinary conditionalization of (ii). The shift to K_2 must, perforce, be a conditionalization on the distribution (ii) conditionalized already on e .

However, if the shift had been direct from K to K^* , the shift would have been via conditionalization on (i) and not on (ii).

Thus, it is entertainable that D -conditionalization will require endorsing different credal states relative to the same corpus.

Equivocation can be avoided by requiring that credal states be the largest convex sets compatible with an interval-valued specification of degrees of credence.

Dempster does, in point of fact, follow Smith and Kyburg in adopting such a restriction on credal states although he fails to explain why he does so in anything like the manner just indicated.† Indeed, in the papers on which I am relying here,

† See "Upper and Lower Probabilities," pp. 330-331 and "A Generalization of Bayesian Inference," p. 245. Dempster's classification of views is somewhat misleading. He recognizes three views: Those which allow any convex set of Q -distributions to be a credal state (the set of such convex sets being Ω); those which allow any convex sets belonging to Ω , consisting solely of convex sets "defined solely by inequalities on probabilities of events"; and those which allow any set belonging to Ω_2 of such sets satisfying Dempster's requirement of representability by a multivalued mapping from a space over which an n -function is defined. Ω_1 appears to be the set of all convex sets representable as the largest convex sets consistent with interval-valued specifications. Dempster claims that Smith allows any member of Ω . My reading of Smith's "Consistency in Statistical Inference and Decision" (*J. Royal Stat. Soc.*, ser. B, v. 23 (1961), pp. 1-37) suggests that he restricts credal states to convex sets in Ω_1 as does Kyburg. The evidence is that Smith defines upper and lower probabilities in terms of upper and lower odds on propositions. Dempster correctly recognizes that Smith's procedures determine supporting lines for convex sets. I surmise he thought that Smith extended his methods in something like the way I propose in section 9.11 and, hence, that Smith is committed to allowing all convex sets in Ω to represent credal states.

This is sheer conjecture; for there is no explanation in Dempster of why he located Smith as he does. In any case, Dempster's interpretation of Smith is not, in my view, supported by Smith's paper. Aside from Dempster himself,

Dempster actually imposes a stronger restriction on credal states than that mandated by intervalists.

Consider a space of simple joint propositions consistent with K of the form $e_x \& h_\theta$. Consider also a space of values r together with a multivalued mapping $t(r)$ from values of r onto subsets of the set of simple joint propositions. Let T be a subset of the set of simple joint propositions and t_* the set of all those values for r such that $t(r)$ is equal to T or some nonempty subset of T . t^* is the set of all those r 's such that the intersection of $t(r)$ and T is nonempty.

Finally, suppose we have a probability distribution over the values of r .

The upper probability for the hypothesis that some element of T is true is $n(t^*)$. The lower probability for that hypothesis is $n(t_*)$.

Because there may be a value of r such that $n(r) > 0$ when $t(r) = \emptyset$, to obtain suitably normalized measures for upper and lower probabilities, the values just defined are divided by 1 minus the total n -value assigned to all values of r for which $t(r)$ is nonempty.

Dempster's restriction is that a credal state over the space of simple joint propositions be the largest convex set compatible with upper and lower probabilities characterizable by means of such a multivalued mapping from a system of values r and a distribution over such values.†

Dempster's restriction is not easy to motivate. To understand its ramifications, a simple example is worth exploring.

Let the space of simple joint propositions consist of $e_H \& h_{.6}$, $e_H \& h_{.4}$, $e_T \& h_{.6}$, and $e_T \& h_{.4}$, where the $h_{.6}$ and $h_{.4}$ specify the chance of heads on a toss of coin a to be .6 and .4, respectively, and e_H and e_T report the results of a toss to be landing heads and landing tails, respectively.

I am, to my knowledge, the first to suggest allowing any member of Ω to be a credal state; and Dempster, in point of fact, merely entertained the idea in passing without in any way endorsing it.

† Convex sets in Ω_2 of the preceding footnote are the largest convex sets satisfying the interval-valued specifications obtained in the manner just indicated in the text. Dempster explains these restrictions in both "Upper and Lower Probabilities" and in "A Generalization of Bayesian Inference," pp. 208-209.

There are 16 subsets of the set of simple joint propositions. Let us consider the first 16 positive integers as the values of r and let $t(r)$ be defined as follows:

- $t(1) = \emptyset$
- $t(2) = \{e_H \& h_{.6}\}$
- $t(3) = \{e_H \& h_{.4}\}$
- $t(4) = \{e_T \& h_{.6}\}$
- $t(5) = \{e_T \& h_{.4}\}$
- $t(6) = t(2) \cup t(5) \quad (d_1)$
- $t(7) = t(3) \cup t(4) \quad (d_0)$
- $t(8) = t(2) \cup t(3) \quad (e_H)$
- $t(9) = t(4) \cup t(5) \quad (e_T)$
- $t(10) = t(2) \cup t(4) \quad (h_{.6})$
- $t(11) = t(3) \cup t(5) \quad (h_{.4})$
- $t(12) = \overline{t(2)}$
- $t(13) = \overline{t(3)}$
- $t(14) = \overline{t(4)}$
- $t(15) = \overline{t(5)}$

$t(16) =$ the entire space of simple joint propositions.

Dempster's restriction is that a credal state over elements of the space of simple joint propositions is the largest convex set determined by specifications of upper and lower probabilities for all hypotheses generated by the four simple joint propositions determined in the manner indicated before by a probability distribution over the values of r .

Suppose, as a good objectivist necessitarian would, that the lower probability for $h_{.6}$ is 0 and the upper is 1 (so that the same interval is assigned to $h_{.4}$ as well).

The set T associated with $h_{.6}$ is $t(10)$. Hence $t_* = \{2, 4, 10\}$. Because the lower probability for $h_{.6}$ is 0, $n(t_*)$ should equal 0. Hence, $n(\{2, 4, 10\}) = 0$. Hence, $n(\{2\}) = n(\{4\}) = n(\{10\}) = 0$.

By similar reasoning, we have $n(\{3\}) = n(\{5\}) = n(\{11\}) = 0$.

A good objectivist necessitarian would assign the pivotal hypotheses d_1 and d_0 the Q -values .6 and .4, respectively, via

direct inference. Consider d_1 . The set T is $t(6)$, $t_* = \{2, 5, 6\}$, and

$$t^* = \{2, 5, 6, 8, 9, 10, 11, 12, 13, 14, 15, 16\}.$$

Because the upper and lower probability is .6, $n(t_*) = n(t^*) = .6$ (provided we assume $n(1) = 0$. This assumption does no harm.) Hence, the elements in the difference between t^* and t_* must bear 0 n -value. Hence only $n(\{2\})$ and/or $n(\{5\})$ and/or $n(\{6\})$ can be positive. But we already know that $n(\{2\}) = n(\{5\}) = 0$. Hence $n(\{6\}) = .6$. By similar reasoning $n(\{7\}) = .4$. Since these n -values sum to 1, all other values of r bear 0 n -value.

From this it is easy to establish via Dempster's restriction that the lower probability for e_H must be 0 and the upper probability must be 1, and similarly for e_T .

An objectivist necessitarian would, however, assign to e_H a lower probability of .4 and an upper probability of .6.

What has been shown is that if Dempster's requirement on credal states is imposed, three conditions which objectivist necessitarians would impose on the credal state cannot be simultaneously satisfied:

- (i) that the lower probability for $h_{.6}$ be 0 and the upper probability 1;
- (ii) that the probability for $d_1 = .6$;
- (iii) that the lower probability for e_H be .4 and the upper probability be .6.

Perhaps we should not obligate rational agents to endorse the dictates of objectivist necessitarianism in this case. Indeed, I have been arguing that we should not do so. But neither should we prohibit joint satisfaction of these three conditions. This, however, is what Dempster's requirements do.

To my knowledge Dempster does not consider precisely this sort of case. But by implication, he seems to favor endorsing (i) and (iii) and rejecting (ii).¹³ Calculation using Dempster's restriction leads to $n(\{8\}) = n(\{9\}) = .4$ and $n(\{16\}) = .2$. The lower probability for d_1 becomes 0 and the upper probability 1. This, of course, is in flagrant violation of the principle of direct inference—not only according to my principle but according to Kyburg's and to everyone else's.

This consideration suggests that Dempster's restriction on credal states is excessive, and that he should have rested content with endorsing the principle of D -conditionalization and the requirement that credal states be representable by the largest convex sets consonant with the upper and lower probability specifications. However, matters are not so simple.

The principle of direct inference as I formulated it in chapter 12 applies to a corpus K containing knowledge of a simple chance hypothesis and knowledge of stochastic irrelevance concerning information about a suitable kind of trial. To obtain counsel about direct inference when the knowledge of chances is composite, confirmational conditionalization had to be invoked.

Consequently, if Dempster is construed as endorsing direct inference (as I formulate it) and D -conditionalization, (ii) must be endorsed; but even if (i) is adopted, it does not follow that (iii) is implied as well.

The principle of direct inference mandates that the expansion of K obtained by adding $h_{.6}$ (call it $K_{h_{.6}}$) is such that $C(K_{h_{.6}})$ mandates .6 as the Q -value for e_H . From this, however, it does not follow that $Q(e_H; h_{.6}) = .6$ for every Q -function in $C(K)$ as it would were confirmational conditionalization endorsed. What does follow is that $Q(e_H; h_{.6}) = .6$ for each Q -function in $C(K)$ for which $Q(h_{.6})$ bears its maximum value. If condition (i) is met, that upper probability is 1.

From this we obtain the result that the upper probability of e_H & $h_{.6}$ must be .6—provided, of course, that condition (ii) is met, as it must be if direct inference is satisfied. On the other hand, if condition (i) is satisfied, e_H & $h_{.6}$ can bear a Q -value of 0. Hence, the interval range is from 0 to .6 as it is for e_T & $h_{.4}$.

Continuing the calculations, we see that the lower probability for e_H & $h_{.4}$ and e_T & $h_{.6}$ remains 0 while the upper is .4.

Direct inference, coherence, and D -conditionalization do not impose further restrictions. Hence, there is nothing preventing the assigning of the total Q -value to e_H & $h_{.4}$ and e_H & $h_{.6}$ so that the upper probability of e_H is 1 for a necessitarian who endorses objectivism but shifts to D -conditionalization.

The upshot is that such a Dempsterian objectivist necessitarian would adopt (i) and (ii) but abandon (iii). He would,

thereby, conform to the restrictions on credal states imposed by Dempster.

Consequently, a modified objectivist necessitarian who changes that doctrine by shifting from confirmational conditionalization to *D*-conditionalization cannot escape the difficulties which Dempster's own restrictions on credal states engender.

I do not know whether Dempster developed his approach as a response to the problems facing objectivist necessitarianism—although his discussion does suggest that he did. It seems clear that he did not follow Dempsterian objectivist necessitarianism as I have reconstructed it. Hence, I cannot pretend to have explained his own motivation for his work.

However, it is true of his scheme as he presents it that the coin example cannot be analyzed in a manner satisfying (i), (ii), and (iii). It is also true that *D*-conditionalization, a principle he endorses, implies this untoward result when combined with objectivist inductive logic.

It seems to me that even if one began with an open mind concerning the viability of modifying confirmational conditionalization, these results would be unsatisfying and reinforce the suspicion that confirmational conditionalization ought not to be abandoned.

In my view, violation of confirmational conditionalization is sufficient reason to reject an approach to probabilistic judgment. My basic case is predicated on the considerations introduced in chapter 10. Such considerations are reinforced by the untoward consequences which tampering with confirmational conditionalization incurs. This is illustrated by our consideration of *D*-conditionalization.

16.6 Random Selection vs. Random Membership

The principle of direct inference, so I have maintained, is required as a way of elucidating the links between knowledge of chance and judgments about test behavior. Such links are not needed if one dispenses with chance predicates altogether as B. De Finetti advocates doing. But if one dispenses with chance, the explanatory and predictive benefits which information about chances provides through the operation of the principle of direct inference will be lost.

An important feature of Kyburg's theory is that it allows him to have De Finetti's cake and to eat it too. Kyburg's

inductive logic does not obligate him to eschew chance; but, in point of fact, Kyburg agrees with De Finetti in regarding chance as a form of *Unsinn*.

Nonetheless, Kyburg claims that direct inferences may be made from knowledge of the truth of frequency statements. Knowledge of relative frequencies provides the explanatory and predictive benefits which knowledge of chances is supposed to furnish.

In illustrating Kyburg's approach to direct inference, extensive use was made of a situation where selections were made from the population of Swedes. In such examples, it may appear that knowledge of chances is expendable. After all, appeal was made to the percentage of Protestants among the Swedes, among the visitors to Lourdes and among Swedish visitors to Lourdes.

Appearances are misleading; for *X* was required to know that the individual was obtained on a trial of kind *S*—i.e., on a random selection from the population consisting of the union of the members of the group who have visited Lourdes and those who are Swedes.

To say that the selection is random is to say that the chance of obtaining one member of the population is equal to the chance of obtaining another. An implication is that the chance of obtaining a Protestant on a trial yielding a Swede is equal to the percentage of Protestants among the Swedes. Hence, given that the selection is a random one, there is a connection between chances and frequencies in the population. But to justify a direct inference (so I maintain) one must have knowledge of the connection given through knowledge that the mode of selection from the population is random.

Kyburg's analysis of such situations is quite different. I ignored the difference in my original presentation of his view. This did not distort the discussion; for his account of direct inference is applicable both to knowledge of chances and to knowledge of frequencies in populations.

My theory, on the other hand, is applicable only for direct inference from knowledge of chances. There can be no direct inference of any value unless one has knowledge of chances.

Kyburg's theory does allow for direct inference from knowledge of relative frequencies whether there are chances or not.

If I am right in rejecting his theory because it violates

confirmational conditionalization, a strong case can be made that one cannot follow Kyburg in rejecting the concept of chance as nonsense and yet provide for direct inference. One must either take De Finetti's view that chance is nonsense—and, hence, that there is no need for direct inference—or acknowledge the significance of chance.

I cannot claim proof that these alternatives are exhaustive if Kyburg's theory is rejected. There may be variants of Kyburg's theory which avoid violating confirmational conditionalization; but I do not know of any worth pursuing.

In spite of the affinity of Kyburg's view to frequency interpretations of chance or statistical probability, he does not endorse such an interpretation. Kyburg simply does not allow conceptions of chance in his conceptual framework. He is entirely right in avoiding the widespread confusion of chance with relative frequency.

The terms R and S which appear in "The chance of an R on a trial of kind S on setup a equals p " occur nonextensionally. They appear extensionally in "100 p % of S 's are R 's." One cannot specify that the truth of the second statement is necessary and sufficient for the truth of the first.

Kyburg does acknowledge that there are statements of probability which are *prima facie* statements of chance or statistical probability. But he takes them to be epistemic statements asserting knowledge of frequencies. Hence, he can account for the nonextensional contexts by supposing them to be located in expressions of propositional attitude. Chance statements are not, therefore, statements of the objective features of chance setups but are subjective or epistemological. The only objective statements involved concern relative frequency, and such statements serve as major premises in the "statistical syllogisms" involved in direct inference. There are no intelligible statements of chance in the sense I favor. To this extent, Kyburg agrees with De Finetti.

According to Kyburg's approach to the Petersen example, X knows that Petersen is a member of the set of Swedes and knows that 90% of Swedes are Protestants. Even if he knows nothing about how Petersen was selected for presentation to him from among the Swedes, X may be entitled to equate the value of his degree of credence for the hypothesis e that Petersen is a Protestant with the value .9.

To mandate this, X 's corpus must be such that any further knowledge of Petersen's membership in various sets must be K -irrelevant in a sense which is analogous to the case of direct inference from knowledge of chances. Thus, in case 8 of section 16.2, X 's knowledge that Petersen is a member of the set of visitors to Lourdes is K -irrelevant and the degree of credence for the hypothesis that he is a Protestant should be .9 even if the example is modified so that it does not concern the results of random selection but concerns knowledge of membership in various sets.

Consider, however, that X knows by deductive closure that Petersen is a member of the set whose only member is Petersen. Is the information K -irrelevant? It clearly is; for all X knows about the percentage of Protestants in this set is that it is either 0% or 100%.

Suppose we attempted to impose an analogue of L -irrelevance in this case. X would have been obliged to consider the information that Petersen belonged to the unit set of which he is the sole member as L -relevant and would, if he were necessitarian, be obliged to assign the degree of credence $[0, 1]$ to the hypothesis that Petersen is a Protestant.

If direct inference is from knowledge of frequencies in sets and knowledge of set membership as Kyburg proposes, imposing the requirement of L -irrelevance always leads to the result that a maximally indeterminate interval will be mandated. (I am, of course, assuming necessitarianism.)

From the point of view of one who, like myself, favors conforming to confirmational conditionalization and imposing L -irrelevance on direct inference, no direct inference of this sort can be worthwhile. This constitutes a decisive objection to Kyburg's program for replacing knowledge of chances in direct inference by knowledge of relative frequencies.

By shifting from the condition of L -irrelevance to K -irrelevance, Kyburg is able to obviate this objection. Of course, he avoids one problem only to confront another—namely, his violation of confirmational conditionalization.

It may, perhaps, be thought that the difficulty involved in using L -irrelevance rather than K -irrelevance in direct inference from knowledge of frequencies applies to direct inference from knowledge of chances as well. This is not so.

Suppose X knows that Petersen is randomly selected from

the Swedes. He also knows that the trial (of kind *S*) has yielded Petersen. What is the chance of obtaining a Protestant on a trial of kind *S* which yields Petersen?

Nothing in logic implies that that chance should be 0 or 1. It could equal the chance of obtaining a Protestant on a trial of kind *S*. In other words, the information that the trial is one which yields Petersen could be stochastically irrelevant. If *X* knows this to be the case, the condition of *L*-irrelevance is met.

Thus, when direct inference is from knowledge of chances, useful direct inference is not automatically precluded by the principle I have proposed.

Perhaps, it will be argued that the chance of obtaining a Protestant on a trial of kind *S* yielding Petersen must be 0 if he is not a Protestant and 1 if he is. The argument is fallacious.

Suppose that designators are assigned to Swedes at random. If *r* out of a total of *n* Swedes are named Petersen, the chance of any subset of *r* Swedes being so designated is the same as any other. If exactly one Swede is named Petersen, the chance of that Swede being so called is $1/n$.

Under these circumstances, the chance of obtaining a Protestant on a trial of kind *S* yielding Petersen is equal to the chance of obtaining a Protestant on a trial of kind *S* and this value can be different from 0 and 1.

Of course, *X* will not always know that the designator or description he uses to individuate the person sampled from the Swedes is a random designator in the sense just indicated. The information conveyed by the designator could be *L*-relevant. In that event, the condition of *L*-irrelevance prevents assigning .9 as the degree of credence to the hypothesis that Petersen is a Protestant.

This is not mere pedantry. Much survey sampling involves using social security numbers as designators for members of a population selected at random. But social security numbers are not random designators. (Suppose, for example, one is concerned with age distribution in a population.) It is not always known that the information about social security numbers is even approximately stochastically irrelevant.

The important moral of the story is that direct inference, as I construe it, depends on knowledge of stochastic irrelevance. Direct inference implies not merely some knowledge of

chances but a substantial amount. One must not only know that the selection of a Swede is random but that his being Petersen is stochastically irrelevant and that the selection being at time *t* is stochastically irrelevant, etc.

Kyburg has complained of this ramification of my view of direct inference.¹⁴ Even if one is concerned to test a statistical hypothesis on the basis of experiment, knowledge of other chances must be presupposed. Consequently, there is no way in which one might reconstruct how knowledge of chances is acquired beginning exclusively with knowledge of data. Kyburg finds this objectionable and is willing to go to great lengths to avoid it. He claims that by using his principle of direct inference one can avoid this result. Because of this, he is prepared to jettison confirmational conditionalization.

My own view is that Kyburg has his priorities wrong. He is still bound to the prejudices of pedigree theories of knowledge which were mentioned at the beginning of chapter 1. In practice, we always find ourselves with a commitment to a substantial amount of background knowledge on which we base judgments of serious possibility and relative to which appraisals with respect to credal probability are made. This background contains assumptions about chance and other theoretical magnitudes as well as information acquired through the testimony of the senses and other more or less reliable witnesses. The origins of this information are quite irrelevant to the deliberations of the agent who has the information.

If that is so, I fail to see the force of Kyburg's objection. What does matter is how judgments of credal probability are grounded on what is in the corpus and how both the corpus and the confirmational commitment ought to be revised over time.

Thus, the issue of whether one should obey confirmational conditionalization or not is of rather greater moment than whether one could reconstruct one's knowledge at a time so that it is all derivable somehow from the testimony of the senses and the records of the memory within one's favorite conceptual framework.

There is, in my view, no important philosophical cost incurred by offering a negative answer to the second question once one is liberated from epistemologies dominated by a

preoccupation with pedigree; but there is substantial cost in giving up confirmational conditionalization.

For me, at any rate, the verdict is clear. Kyburg is to be credited with having focused the issue of direct inference more sharply than any of his predecessors managed to do. He has helped the rest of us to think a little more clearly. But his own positive response to the challenge confronting objectivist necessitarianism should be rejected precisely because its liabilities do outweigh its benefits.

17.1
Observation
Reports

X witnesses the toss of a coin and observes that it lands heads up. His observing that the coin lands heads up is a response to sensory stimulation constrained by his past conditioning and his biological and psychological constitution. *X*'s response shall be called his "making an observation report that the coin lands heads."

Reporting that the coin lands heads is an event just as making the statement that the coin lands heads is. However, *X* may make an observation report without uttering or in any other way making a statement. *X* might acquire the disposition to respond affirmatively to the question "Is *h* true?" in reporting that *h*. Perhaps, his reporting that *h* will be accompanied by some feeling of conviction. I do not care to speculate about such matters. I assume only that men do make observation reports in response to sensory stimulation, that such reportings are episodic propositional attitudes, and that the reportings are conditioned by the agent's biological and psychological constitution as well as any prior conditioning of that agent.

Reporting that *h* is not to be confused with adding *h* to one's standard for serious possibility or corpus. The difference should become apparent when the agent *X* is not an individual person but some group. The testimony of the senses will consist of the reportings of some person or persons who may or may not be members of the group and whose reports may or may not be routinely added to the group corpus. But even when *X* is a person and makes observation reports, we should distinguish between the making of the report that *h* in response to sensory stimulation and the accepting of *h* into evidence for use in subsequent inquiry.

When *X* adds an observation report to his corpus of knowledge, he is not justified in doing so on the basis of some inference from evidence already in the corpus. The expansion

should be regarded as the output of a program X adopted beforehand designed for obtaining new error-free information, where the input is the observation report that h made in response to sensory stimulation. Neither the observation report that h nor the report that X reports that h (that h appears true to X) is a premise of an inference to the conclusion that h is true or that h should be added to X 's corpus.

For example, X may be committed to a program prescribing the addition of any sentence h describing the color of objects to X 's corpus when X reports that h is true. When implementing such a program, X makes no inference from premises to conclusion but lets the report he makes in response to sensory stimulation determine what he shall add to his corpus in accordance with some previously adopted procedure.

In chapter 2, where this idea of routine expansion was introduced, implementing such a program was taken to be equivalent to conducting a trial on a chance setup where the space of possible outcomes consists of possible expansions of X 's corpus. Whereas X is not obliged to justify the expansion resulting from implementing his program by appeal to an argument from premises used as evidence, he may sometimes be obliged to justify his using the program he adopts. From X 's point of view, a program is a good one if it is reliable, and he is entitled to use it if he knows this to be so—i.e., if he knows that the chance of his importing error into his corpus by using the program is sufficiently low.

Insofar as inference is involved in observation, it becomes relevant when the programs X uses are being subjected to critical review. When that happens, X may be called upon to reach conclusions via inferential expansion concerning the reliability of various programs. But the actual implementation involves no inference at all. It is a matter of skill.

Thus, a distinction should be made between (i) making an observation report that h , (ii) using the making of such a report as input into a program for expansion via observation, and (iii) using h as evidence in a corpus of knowledge in subsequent inquiry and deliberation.

It may appear that although implementing a program for expansion involves no deliberation or inference, prior to every occasion on which a program is to be implemented, X is faced with a choice between rival programs for expansion and should

weigh these alternatives relative to what X knows at that time. This is clearly not the usual practice and I do not mean to recommend it.

It is true that prior to making an observation report X must be committed to some program for using that report as input into a program for expansion. If he lacks such a program, he will lose the opportunity to expand via the testimony of his senses. Sometimes X may be in a position to anticipate his future opportunities for making observations and to plan for them. Yet, such opportunities usually occur unexpectedly; and even when they are anticipated, it is often too costly to design a program for each occasion.

In order to exploit the testimony of the senses effectively as a source of reliable information, X should adopt a program applicable on a wide range of occasions easily identifiable so that on each occasion X is in a position to implement the program in a habitual or routine manner. Thus, expansion via the testimony of the senses is typically routine expansion as I claimed in chapter 2.

In sum, expansion via the testimony of the senses has two important characteristics: (a) It is the result of using a program for using observation reports made in response to sensory stimulation as inputs; and (b) it is routinizable, so that X is prepared to implement it to obtain information otherwise inaccessible or too costly.

None of this implies that if X is unprepared beforehand to implement a program for routine expansion, he forgoes all opportunity to add the sentence h reported true to his corpus. He loses only the opportunity to add h on that specific occasion in response to his reporting that h is true.

But X may have other opportunities to make observations. Or X might use reports of the testimony of others as inputs. Finally, X might have sufficient knowledge to warrant an inferential or inductive expansion leading to his accepting h into his corpus.

Nonetheless, X will often lack the knowledge to settle the question of the truth of h via inferential expansion; and the alternative method of routine expansion via the testimony of others may not be sufficiently reliable. At any rate, the testimony of the senses is an important source of information and it is desirable that it be exploited.

In chapter 2, routine expansion, whether via the testimony of the senses or the testimony of witnesses, was distinguished from deliberate or inferential expansion by its capacity to inject contradiction into a corpus. In inferential expansion, *X* evaluates rival potential answers to a given question in the light of the total knowledge available to him in order to obtain error-free information. He should not deliberately expand into contradiction.

If prior to each occasion of observation *X* were in a position to design a program for expansion using observation reports as inputs, his concern to avoid error should preclude his endorsing a program which would lead to his injecting inconsistency into his corpus. But *X* is not in a position to hand-tool a program for expansion via observation on each occasion. Such a procedure is either too costly or not feasible. *X* will seek a program for expansion yielding information—say about colors—which is relatively insensitive to variations in his prior knowledge of the color of objects to be observed on various occasions. Thus, *X* will be in a position to employ an easily routinizable program not requiring close attention to the circumstances (in particular his state of knowledge) peculiar to each occasion of use. *X* will seek a program he can apply in a routine manner whether he knows that the color of the object to be observed is not purple or lacks such information prior to observation.

Thus, *X* might know that the chance of his making a false report of the color of objects under conditions *C* is extremely low but not 0. In a situation where he already knows the color of an object not to be purple, he may continue to assume the chances of making an error very low.

To be sure, *X* could design a routine expressly for that occasion—a routine that recommends adding the report made to the corpus except when the report that the color is purple is made. But the benefit to be obtained thereby is very small if the chance of reporting purple given that the color is purple is already very low. Furthermore, the cost of designing a special program for a special occasion in this manner may serve as a deterrent. Finally, since the chance of reporting purple is assumed low to begin with, should *X*, nonetheless, report that the color is purple, there would be some pressure to contract the corpus to give rival hypotheses a hearing any-

how—just as there would be should the testimony of the senses inject contradiction into the corpus. Under the circumstances, the benefits of routinization outweigh the costs—including the small risk (if the routine is reliable) of injecting contradiction into the corpus.

Writers with an empiricist bias often emphasize the capacity of the observation to contradict our most cherished assumptions. On the view I am developing, this is not a virtue of routine expansion but a cost which is compensated for by the circumstance that routine expansion via observation furnishes a new information in an efficient and fairly reliable manner.

Thus, the kernel of truth in empiricism is the contention that some routines for expansion via observation reports are both highly reliable and yield useful information. To be sure, in the course of inquiry, our judgments of the reliability of such routines (like other knowledge) is open to revision and so may our judgments of the fecundity of routines in yielding new information. We cannot claim very *fixed* knowledge of which routines are reliable.

The use of observation reports in routinizable programs is not restricted to routines for expansion. Programs can be devised and implemented where the outputs are actions of various kinds. As in the case of routine expansion, observation reports are used as inputs without being used as part of the evidence on the basis of which the act performed is identified as admissible from among the available alternatives.

17.2 Objectivist Ne- cessitarianism a la Neyman- Pearson

Objectivist necessitarianism implies that observation reports admitted into evidence will, in general, be confirmationally irrelevant for statistical hypotheses being subjected to test. Exceptions occur when such observation reports entail the hypotheses under test; but this is of small consolation when dealing with statistical hypotheses.

Objectivist necessitarianism seems to imply, therefore, that admitting observation reports into evidence is without value in much subsequent inquiry and deliberation. The presumption that the data of observation are of value in scientific inquiry conflicts with this apparent implication and constitutes the core of the objection raised previously against objectivist necessitarian doctrine.

The previous three chapters canvassed ways and means of

modifying objectivist necessitarianism so as to secure confirmational relevance for observation reports without straying too far from the spirit of objectivist necessitarian doctrine. The considerations adduced suggest that none of these approaches are satisfactory.

My conclusion is that necessitarianism should be abandoned in favor of a revisionist outlook. But there is one other important alternative remaining to be considered. Perhaps we have been too hasty in concluding that the data of observation are without value according to objectivist necessitarianism.

There is no doubt that the implication that in many important contexts such data are confirmationally irrelevant is valid. What remains to be considered is whether that implication has the further consequence that the data are without value.

What has been shown is that the sentences reported true in making observation reports are useless *when accepted as evidence and used in this capacity*. What remains an open question, however, is whether the making of observation reports is not useful *as input* into programs for selecting actions, policies, or expansion strategies.

In 1933, J. Neyman and E. S. Pearson granted that using observation reports as evidence is useless:

We are inclined to think that as far as a particular hypothesis is concerned, no test based on the theory of probability can by itself provide any valuable evidence of the truth or falsehood of that hypothesis.¹

Neyman and Pearson go on to say in a footnote that exceptions arise when observation reports entail the falsity of a hypothesis.

I. Hacking misinterprets these remarks when he attributes to Neyman and Pearson the view that "there is no alternative to certainty or ignorance."² Neyman and Pearson were prepared to recognize numerically precise credal states as legitimate, provided they could be derived from knowledge of chances via direct inference. Like R. A. Fisher, they were ready to endorse the use of Bayes' theorem to derive posteriors from priors provided the priors could be obtained through direct inference from knowledge of chances.

As I read them, Neyman and Pearson concede that when priors cannot be derived via direct inference, the use of ob-

servation reports to obtain posteriors via Bayes' theorem is foreclosed. That is why observation reports fail to provide "valuable evidence of the truth or falsehood of that hypothesis." This reading is supported by Pearson in 1962:

We [i.e., Neyman and Pearson] were certainly aware that inferences must make use of prior information and that decisions must take account of utilities, but after some considerable thought and discussion round these matters we came to the conclusion, rightly or wrongly, that it was so rarely possible to give sure numerical values to these entities, that our line of approach must proceed otherwise. Thus we came down on the side of using only probability measures which could be related to relative frequency.³

Hacking himself points out that Neyman and Pearson are committed to the legitimacy of inferring from knowledge that in a long series of stochastically independent trials in which a given kind of outcome occurs with constant chance p that it is "practically certain" that the relative frequency with which that kind of outcome occurs approximates p . He thinks he detects an inconsistency between this view and the thesis that "there is no alternative to certainty or ignorance."† There is, indeed, an inconsistency between these two views; but only one of them is attributable to Neyman and Pearson.

The very circumstance that Hacking's reading of the Neyman-Pearson doctrine renders it inconsistent when there is a readily available alternative interpretation argues against his construal.

According to my interpretation of the passages cited, Neyman and Pearson are affirming the implications of objectivist necessitarianism. They are prepared to accept the conclusion

† Hacking (*Logic of Statistical Inference*, Cambridge: Cambridge University Press, 1965, pp. 104-105) makes much out of a claim by Neyman and Pearson which implies that it can be "proved" that the relative frequency approximates p . He correctly notes that no proof in logic and probability theory can be given. But a more charitable interpretation would allow Neyman and Pearson "proofs" which rely at least implicitly on a principle of direct inference. Hacking himself claims, somewhat glibly I think, that there is a "principle about support and chance" which "seems universally to be accepted that it is hardly ever stated" (*ibid.*, p. 135). This principle is what Hacking calls the "frequency principle" and what I call "direct inference." Hacking fails to recognize the presence of substantial disagreements in the literature concerning the content of such a principle of the sort discussed in the previous chapter. But insofar as the principle is universally accepted or nearly so, why does Hacking think that Neyman and Pearson are exceptions?

that observation reports are confirmationally irrelevant when used as evidence except, of course, when they entail the falsity of some hypotheses subject to test.

Having bitten the bullet, however, they are not prepared to rest content with the uselessness of observational data in inquiry and deliberation. Following immediately on the heels of the passages cited from the 1933 paper, they write:

But we may look at the purpose of the tests from another view-point. Without hoping to know whether each separate hypothesis is true or false, we may search for rules to govern our behavior with regard to them, in following which we insure that, in the long run of experience, we shall not be too often wrong.⁴

What Neyman and Pearson mean by "rules to govern our behavior" is not entirely clear. I suggest, however, that a good approximation to their intent is obtained by construing them as recommending the use of programs for using observation reports as inputs into programs designed to select acts. Moreover, they appear to favor routinizable programs which may be used over and over again in appropriately similar situations and where the chance of success is high on each occasion so that in the long run it is "practically certain" that "we shall not be too often wrong."

A familiar illustration of the sort of situation where Neyman-Pearson methods are alleged to be appropriate is in the design of procedures for quality inspection of goods at an appropriate state in the manufacturing process. The plant manager wishes to provide instructions to quality inspectors which they are to adhere to in a regular manner when passing on the suitability of putting the finished article on the market or returning it for reprocessing. By routinizing the inspection procedure, costs of inspection are reduced, reliance on the judgment of the inspectors is reduced, while the reliability of their verdicts over the long run is secured. From the plant manager's point of view, the observation reports made by the inspectors are used as inputs into a routinizable program for checking on the quality of the finished product. (The alienation of the quality inspector from his labor is even more extreme than what most Marxists dream of.) In this example, the plant manager uses the observation reports of others as inputs into programs. We

may entertain, however, using our own reports as inputs into such programs as well.

Perhaps objectivist necessitarianism is untenable solely because it implies the confirmational irrelevance of data and, hence, the uselessness of observation reports as evidence. But if we are going to give objectivist necessitarianism a run for its money, we should explore the extent to which observation reports may be usefully exploited as inputs in programs.

The question acquires additional interest from the circumstance that, if my interpretation of Neyman-Pearson doctrine is sound, the dominant outlook among students of contemporary statistical theory endorses the view to be examined.

That the making of observation reports can legitimately serve as inputs in programs for expansion or performing other acts is undeniable. To save objectivist necessitarianism, however, what must be shown is that in important contexts where objectivist necessitarians concede that observation reports are useless as evidence but should, according to presystematic judgment, have value, they do have value as inputs into programs.

My main contention is that this thesis can be defended only by legislating strong conditions on the values and goals men should seek to promote. In the chapters immediately preceding this one, I objected to efforts to impose conditions on rational credence stronger than those mandated by objectivist inductive logic and to claim that they are principles of inductive logic obligatory on all rational agents under all circumstances. I did not object to imposing some of these conditions under special circumstances. By the same token, I take exception to views which claim that rational men should endorse constraints on goals and values stronger than the fairly weak ones advanced here in order to qualify as rational. I do not object to imposing such stronger constraints under special circumstances. Indeed, X may very well be committed to some special system of values and goals entailing much stronger restrictions on his valuations of hypotheses about consequences than the conditions on rational valuation imply. But agents should be allowed to conform to the dictates of principles of rational deliberation and valuation even when they differ widely in the goals and values they endorse. Perhaps, one party or the other (or both) in some dispute over values

is immoral, imprudent, impolitic, or the like. But because his values are offensive, he is not, therefore, irrational.

Objectivist necessitarianism has to abandon the relative insensitivity of principles of rational choice to variations in goals and values if it is to rescue itself from difficulty along the lines just sketched. In my opinion, this is untenable.

Thus, adopting routine programs for using observation reports as inputs in selecting policies rather than using such reports as evidence in deliberate decision making may be desirable relative to the agent's goals and values. He might be interested in long-run benefits from repeated application of the program and in reducing costs of deliberation on each separate occasion. In such cases, objectivist necessitarians may be in a position to recognize making observation reports as having a value for the decision maker.

This does not entitle objectivist necessitarians to insist that rational agents should always assess benefits in terms of the long run and should evaluate costs in such a manner as to favor routinization over deliberation.

Perhaps, however, objectivist necessitarians can show that even when routinization is not deemed desirable, it is better to employ a program using reports as inputs than to use observation reports as evidence. My contention is that, even in such cases, goals and values would have to be legislated to an excessive degree.

17.3 Security Levels

In deliberation aimed at deciding between rival feasible options in order to realize some system of goals and values (whether "practical" or "theoretical"), objectivists are committed to recognizing as *E*-admissible the largest set of feasible options which are "Bayes solutions" for some probability distribution compatible with the specifications of the problem. Because the set of *E*-admissible options tends to be so large, objectivist necessitarians have sought further criteria for arbitrating between rival *E*-admissible options.

A. Wald's effort to develop a systematic theory of statistical decision making was noted in chapter 7. He claimed that his approach could serve as a generalization on and improvement of the approach of Neyman and Pearson.⁵ Basic to his approach was the use of the criterion, which singles out as admissible members of a subset of the *E*-admissible options.

In chapter 7, I interpreted Wald's criterion as a special case of the principle of *S*-admissibility I have proposed applicable, in particular, in those contexts where credal states are as the objectivist necessitarians recommend.

Not all authors belonging to the Neyman-Pearson "school" follow Wald. However, the main points of criticism applicable to objectivist necessitarianism on the assumption that objectivist necessitarians follow Wald's approach should apply *mutatis mutandis* to variant approaches.

In any case, I have already endorsed the criterion of *S*-admissibility and, as a consequence, exploring the ramifications of objectivist necessitarianism within the framework for rational choice I have been using amounts to exploring the ramifications of maximizing.

In chapter 7, I claimed that the identification of security levels for feasible options depends on how the space of hypotheses concerning the consequences of a given option is partitioned. There are many ways of determining security levels for each feasible option and, hence, diverse ways of determining *S*-admissibility. The methods for fixing security levels which are used in practice belong to a fairly restricted set. Nonetheless, as I contended in chapter 7, there is no principle of rational choice which favors one method over another. The choice of a method of fixing security levels is a feature of the agent's goals and values just as the system of permissible utility functions is.

In chapter 7, I considered two obvious methods for determining security levels which I labeled (a) and (b). I noted that Wald favored method (b).†

† A. Wald, *On the Principles of Statistical Inference*, Notre Dame, Ind.: University of Notre Dame Press, 1942, p. 44. Method (b) is, in effect, determining security levels by means of what Wald calls a "risk function." Sympathetic critics of Wald's theory, such as E. L. Lehman (*Testing Statistical Hypotheses*, New York: Wiley, 1959, pp. 13-14), do not question the propriety of using risk functions in determining security levels. The fundamental complaint is that Wald fails to furnish an account of rational statistical decision making applicable to situations where *X* lacks sufficient prior information to warrant numerically precise priors but, nonetheless, has some information which indiscriminate use of maximin ignores.

In my view, the complaint is just. Some of the proposals made in the present book are designed to meet it (e.g., restricting maximin to the set of *P*-admissible options). However, the Wald theory remains intact in those cases where priors are maximally indeterminate as they would be for objec-

Given a commitment to method (b), it is possible to use observation reports as inputs into programs to great effect even when the programs are not routinized. On the other hand, if security levels are determined by method (a), the use of observation reports as inputs can be no more effective than using observation reports as evidence.

Thus, Wald succeeds in establishing the value of data by adopting a specific method for fixing security levels. But the adoption of such a method is not mandated by principles of rational valuation or rational choice. Perhaps ethical, political, economic, or prudential considerations favor some method or other on some occasion. But it is inappropriate to insist that all rational agents should adopt method (b) so that the viability of objectivist necessitarian doctrine can be saved.

17.4
Mixed
Options

X has an opportunity to accept gamble G_1 where he receives .5 utiles if h_6 (the chance of heads on a toss of coin a is .6) is true and loses .5 utiles if h_4 is true or to accept gamble G_2 where he loses .4 utiles if h_6 is true and wins .4 utiles if h_4 is true. X knows that either h_6 or h_4 is true and is assured that he will be paid off according to the terms of the gamble depending on which hypothesis is correct. K is the initial corpus and X is an objectivist necessitarian.

Case 1: X is compelled to choose between G_1 and G_2 relative to K containing the information just specified and no other pertinent information.

The credal state for h_6 and h_4 is maximally indeterminate. Both options are E -admissible and P -admissible. The security levels for each of the gambles is determined by the values of the gambles contingent on the truth of h_6 and h_4 . There will be no controversy about this. G_2 has a security level of $-.4$, and G_1 of $-.5$. G_2 is uniquely S -admissible and should be chosen.

Case 2: Let X 's options be enlarged in the following way: He is told there is a spinner which, when set in motion, stops with

tivist necessitarians. In my opinion, the serious trouble is that the viability of the Wald theory derives from its requiring the use of minimax relative to risk functions—i.e., relative to method (b) for fixing security levels.

equal chance within one arc on the dial as it does within any other arc of equal arc length. Consequently, the dial can be divided into two portions so that the chance of falling into one portion is p and into the other is $1 - p$. X can give instructions to agent Y as follows: Y is to adjust the apparatus so that p has some particular value between 0 and 1. The spinner is to be set in motion. If Y observes the pointer landing in the first section, he is to select option G_1 and X is to reap the benefits or penalties. Otherwise Y is to select G_2 with X receiving the payoff.

In case 2, X uses Y 's observation reports as inputs into a program for selecting a gamble. X need not find out which gamble is selected until after it is selected. At no point prior to selection does he add or have the opportunity to add Y 's observation report to his body of knowledge to use as evidence. Furthermore, X could have dispensed with a human observer (assuming that the technology is available) and programmed some computer which would print out the gamble to be selected in response to the behavior of the spinner.

To avoid needless complications suppose that X accepts in his corpus of knowledge K that Y is a perfectly reliable observer of the outcome of the spin and will faithfully perform his instructions. The chance of Y making a false observation report is known to be 0. Hence, X knows that if he selects a given value of p , the chance of G_1 being selected is p . By direct inference, X 's degree of credence for the hypothesis that G_1 will be selected conditional on his choosing a specific value of p is equal to the value of p chosen. From X 's point of view, choosing the option A_p (i.e., choosing the value p) is choosing a "mixture" of the pure options G_1 and G_2 .

Mixed options were considered in chapter 7. The new element introduced here is the emphasis on the fact that in choosing a mixed option, X is choosing a program for using data as input. In chapter 7, however, it was already noted that there are at least two methods for assessing security in such cases.

If method (a) is employed and $0 < p < 1$, A_p has four possible consequences: choosing G_1 when h_6 is true or when h_4 is true and choosing G_2 when h_6 is true or when h_4 is true.

The payoffs are .5, $-.5$, $-.4$, .4. The security level is $-.5$ as it is for the option G_1 (when $p = 1$). Only when $p = 0$ so

that the option is equivalent to choosing G_2 is the security level $-.4$. This option is uniquely S -admissible.

If method (b) is employed, A_p has two relevant consequences h_6 and h_4 . The payoffs are the expected utilities conditional on these hypotheses. These are $.5p - .4(1 - p) = .9p - .4$ and $.4(1 - p) - .5p = .4 - .9p$, respectively. If X 's goals and values warrant fixing security levels in this manner, the option with the largest security level is A_p for $p = 4/9$.

Thus, if method (b) is employed for determining security levels, Y 's observation reports concerning the outcome of the toss of the spinner have some value as input into a program for selecting a gamble from X 's point of view. As noted in chapter 7, this is the method Wald favored.

There are four aspects of case 2 which bear emphasis:

(1) X has the opportunity to use someone else's observation reports as inputs into a program for selecting an option.

(2) X 's choice of a program is not routine. Prior to implementation of the program chosen, X considers which of the several programs available to him he should choose in the light of what he knows and his goals and values.

(3) Even if X were in a position to add the observation reports made by Y to his corpus and use them as evidence, the reports would be confirmationally irrelevant to h_6 and h_4 . This confirmational irrelevance of observation reports about the outcome of the trial on the spinner is *not* a consequence of objectivist necessitarianism peculiar to it. Even if X were to endorse a numerically precise prior credal state for h_6 and h_4 , the information he would obtain about the outcome of the spin should not alter his credal state for these two alternatives. Consequently, in case 2, Y 's reports would be useless as evidence regardless of whether objectivist necessitarianism is endorsed or not. Whether they are useful as input into a program for selecting a gamble does not depend, therefore, on their confirmational relevance or irrelevance to the hypotheses of concern.

(4) The usefulness of Y 's reports as inputs does depend, however, on how X fixes his security levels in evaluating S -admissibility. If X uses method (a), the reports are without value. If method (b) is employed, they have value. But X cannot be condemned as irrational for using method (a) rather

than method (b). Which method is employed is a question of goals and values. This does not imply that values are immune to critical review. However, an account of rational choice or goal attainment should restrict itself to imposing weak constraints on goals and values and leave further critical judgment to inquiry into morals, politics, the law, and other domains of practical wisdom.

17.5 Using One's Own Reports as Input

Suppose X uses his own reports concerning the outcome of the spin as inputs into a program for selecting one of the gambles. This predicament resembles case 2 in all relevant respects except the first of the four aspects listed in section 17.4. However, this modification introduces some special complications which warrant distinguishing it from case 2. For this reason we have

Case 3: Exactly as in case 2 except that $Y = X$.

In case 2, X has the opportunity to choose a pure or mixed option before Y witnesses the result of the spin. In case 3, X has the opportunity to choose a pure or mixed option before X himself witnesses the result of the spin. However, in case 3 but not in case 2, X can wait until after he has made the observation report of the outcome of the spin and added this report to his corpus via routine expansion and then decide which of the two gambles to accept on the basis of the information in his expanded corpus.

Situations can, indeed, arise where X is prevented from waiting until he obtains observational evidence in this way and is constrained to make a decision between programs before the spin is made and its outcome observed. In such cases, there is very little of interest distinguishing case 3 from case 2. When, however, X is not constrained in this way, there is something of interest to explore.

Because the reports about the outcome of the spin are confirmationally irrelevant to h_6 and h_4 regardless of whether objectivist necessitarianism is endorsed or not, waiting and using the report as evidence for choosing a gamble must lead to the same decision (choosing G_2) favored in case 1 and in case 2 when method (a) for fixing security levels is adopted.

Conflict arises, however, if X fixes security levels in accordance with method (b). In contemplating what he should

do prior to witnessing the outcome of the spin, X assesses not only the several mixed options available to him but he can also calculate what he should do if he were to base his decision as to which gamble to choose on the evidence available to him after having made the observation report. Relative to what he knows at the time and given his goals and values—including his method of fixing security levels—he may justly reason that he should pick the mixed option for $p = 4/9$ rather than wait and see what the outcome of the spin will be.

X is not, thereby, violating the total knowledge requirement; for that requirement obligates X to employ the total knowledge available to him prior to making a decision.

It may be objected that X does have the opportunity to postpone decision making until he observes the outcome of using the randomizer. Since doing so will preserve several options—in particular, the pure options G_1 and G_2 — X should do so.

This objection is without force. The option of delay in case 3 is equivalent to choosing G_2 . It does not plausibly represent suspending judgment between two or more E -admissible options—at least according to customary ways of ranking options with respect to strength.

There is yet another objection. Suppose X chooses $A_{4/9}$, makes the appropriate observation report, and then considers (on the basis of the knowledge he acquires via routine expansion) whether he should renege on the program previously adopted or not. Relative to the new corpus, he should choose G_2 regardless of the dictates of the program initially adopted.

This objection too is not decisive. In the first place, X may very well implement the program for selecting a gamble simultaneously with implementing the program for routine expansion and, indeed, if he has a choice in the matter, should do so; for he knows that if he reneges he will choose G_2 which is inadmissible from his point of view prior to observing the outcome of the spin. But even if he cannot fully implement the program for selecting a gamble until after having made the observation report and adding it to his corpus, the costs of renegeing may be too prohibitive.

Thus, although some variants of case 3 may turn out to be such that renegeing on the mixed option initially adopted is justified, there will be other cases where that is not so. To the

extent that renegeing is unjustified, the fact that in case 3 X makes his own observation reports where in case 2 Y makes the observation reports turns out to make little difference to the policy X should follow. Furthermore, although it is necessary in cases 2 and 3 that deliberation be conducted prior to making observations of the outcome of the spin, the mere circumstance of temporal priority is not the salient factor. X can, after all, calculate prior to observation how to use reports as evidence after making observations. What is crucial in both cases 2 and 3 is that observation reports are used as inputs rather than as evidence and that the effective use of such reports as inputs depends on how the agent fixes security levels.

17.6 Turning Evidence into Input

Objectivist necessitarians of the Neyman-Pearson-Wald persuasion have joined game theorists of the Von Neumann-Morgenstern persuasion in extolling the virtues of mixed options. However, if all that the previous analysis succeeds in showing is that observation reports can sometimes be used effectively as inputs in programs for implementing mixed strategies, we shall have failed to make any progress whatsoever towards alleviating the difficulties plaguing objectivist necessitarianism. The observation reports used as inputs in cases 2 and 3 are confirmationally irrelevant to $h_{.6}$ and $h_{.4}$ according to everyone, regardless of whether they are objectivist necessitarians or not.

The trouble with objectivist necessitarianism is that it implies the confirmational irrelevance of reports which presystematically would be considered confirmationally relevant to such hypotheses—e.g., reports concerning the outcome of a sequence of tosses of the coin in question. To avoid the implication that such data are useless, objectivist necessitarians may seek to show that they function usefully as inputs into programs for selecting options.

Case 4: Coin a is to be tossed once. The perfectly reliable observer Y is to report the outcome. X has the opportunity prior to observation to instruct Y to select one or the other of the two gambles for him depending on the outcome of the toss. X does not find out, however, what the outcome of the toss is before he is paid off.

The statement e_H asserts that a lands heads on the toss and e_T that it lands tails. e_H is confirmationally relevant in the strong sense to h_6 and h_4 and is confirmationally relevant provided that the prior credal state is not maximally indeterminate (see sections 10.5 and 10.6). Unlike cases 2 and 3, presystematic judgment would agree that data concerning the outcome of the toss has some value in deliberation concerned in choosing between G_1 and G_2 . Hence, the fact that objectivist necessitarians insist that the prior credal state be maximally indeterminate so that the data are confirmationally irrelevant constitutes a genuine difficulty.

This is not the case with mixed options as in cases 2 and 3. Honest men do differ in presystematic judgments concerning the value of data about the outcome of spinning the spinner. Were such data to be useless, that would not present the deep objection to objectivist necessitarianism that the confirmational irrelevance of e_H seems to pose.

Nonetheless, cases 2 and 3 have been helpful; for they show that even data which is prima facie useless can sometimes have value as input into a program for selecting options provided the method for fixing security levels is of the right sort. The next step in the discussion is to extend the idea to case 4, where the data are prima facie useful.

Excluding all mixed programs, there are four programs available to X in case 4:

B_1 prescribes selecting G_1 regardless of Y 's report.

B_2 prescribes selecting G_1 if Y reports that e_H and G_2 if Y reports that e_T .

B_3 prescribes selecting G_2 if Y reports that e_H and G_1 if Y reports that e_T .

B_4 prescribes selecting G_2 regardless of Y 's report.

The payoffs for these four programs are given in table 17.1. The expected payoffs conditional on h_6 and h_4 are given in table 17.2.

From table 17.2 it is apparent that B_3 is not E -admissible. However, each of the other three programs ranks highest with respect to expected utility relative to some permissible assignment of Q -value to h_6 within the range from 0 to 1. Hence, for an objectivist necessitarian, all options except B_3 are E -

Table 17.1

	h_6 & e_H	h_6 & e_T	h_4 & e_H	h_4 & e_T
B_1	.5	.5	-.5	-.5
B_2	.5	-.4	-.5	.4
B_3	-.4	.5	.4	-.5
B_4	-.4	-.4	.4	.4

Table 17.2

	h_6	h_4
B_1	.5	-.5
B_2	.14	.04
B_3	-.04	-.14
B_4	-.4	.4

admissible. They are also P -admissible assuming that no one of them represents suspense between any others.

If method (a) is used to determine security levels, the first table reveals that the security level for B_4 is best and it should be chosen. But B_4 is equivalent to G_2 . Using the data as input is of no value in deciding which gamble to choose.

If method (b) is employed, the second table reveals B_2 to be uniquely S -admissible. Thus, the data of observation are effectively used as inputs in determining which gamble to select.

Suppose X had not been an objectivist necessitarian, but had adopted a prior credal state for h_6 and h_4 assigning to the former alternative credence equal to some subinterval of the interval $[.4, .6]$. If X were to observe the toss or to have Y do so and add the report to his corpus to use as evidence in basing his decision concerning which gamble to accept, he would choose G_1 if he found out that e_H is true and G_2 if he found out that e_T is true.

Thus, X would simulate the program B_2 even though he is deliberately choosing a gamble on the basis of the total evidence to him after observation.

It is a fallacy, however, to point to this simulation and then conclude that an objectivist necessitarian who chooses program B_2 for using data as input is really using the data as evidence and adopting a prior credal state sharper than that prescribed by objectivist necessitarianism. Notice that had X

retained a credal state assigning the interval [.4, .6] to $h_{.6}$ but altered his method of fixing security levels, he would still simulate B_2 . But if X were objectivist necessitarian but had shifted to method (a) for fixing security levels, he would not follow program B_2 .

Case 5: Satisfies the conditions for case 4 except that $Y = X$.

The relation between case 5 and case 4 is analogous to the relation between case 3 and case 2.

There is, however, one point worth emphasizing. In case 3, when choosing between mixed options, it is crucial that X avoid adding the observation report he makes to his corpus until after the program selected is implemented. Everyone, whether they are objectivist necessitarians or not, acknowledges somehow that the observation reports are confirmationally irrelevant in such situations. They can be of use only if they are employed as inputs into programs. Hence, it is desirable to avoid adding them to evidence until a program has been chosen and implemented.

In case 5, however, presystematic judgment suggests that the observation reports are of value in a wide variety of cases and the value does not seem to depend on whether X is entitled to avoid using the information he eventually obtains concerning the outcome of the toss of coin a as evidence in choosing a gamble.

Objectivist necessitarians, however, must insist as a matter of principle that X avoid using such information as evidence.

This appeal to presystematic judgment cannot be considered decisive against objectivist necessitarianism; but the implication is disturbing.

But even if we waive this point, X must fix security levels according to method (b) rather than method (a) in both cases 4 and 5. As I have argued repeatedly, to obligate X as a rational agent to follow this course is to confuse moral, political, or other valuation with conditions on rational valuation.

I have complained previously of the tendency of many authors from Carnap to Hacking to impose as principles of inductive logic strong conditions on credal states and to condemn those who violate these principles as deviants from perfect rationality. In my view, Wald took a step replicating

this practice in connection with questions of value. I do not claim that rational agents should be prohibited from fixing security levels in accordance with method (b); but they should not be coerced or cajoled into doing so merely to save the viability of objectivist necessitarianism.

Thus, although it is undeniable that observation reports may be used to great advantage as inputs into programs for selecting options even when these programs are not routines for expansion, doing so depends for its legitimacy to a considerable degree on the goals and values of the agent. When the goals and values fail to meet the required conditions, we should not condemn the decision maker for his irrationality. It is objectivist necessitarianism that is in trouble.

17.7 The Long Run

The rationalization of the use of observation reports as inputs into programs for action illustrated in cases 2-5 concerns justifying the use of programs in contexts where the agent is interested solely in the gains and losses pertaining to decisions taken in that context.

Followers of Neyman and Pearson, however, often make appeals to the long-run benefits of repeated application of a program. What has been shown is that, provided that the agent fixes security levels in a certain way, such appeals to the long run are unnecessary. On the other hand, it remains worthwhile to consider ways of justifying the use of programs to realize long-run benefits.

Consider first a situation exactly like case 4 or case 5 except that the alternative statistical hypotheses are $h_{.99}$ and $h_{.01}$. Prior to experimentation, X 's credal state is such that the hypotheses e_T & $h_{.99}$ and e_H & $h_{.01}$ have credence representable by the interval [0, .01]. Consequently, if X were to take all four simple joint propositions as an ultimate partition and assign to each element of U an M -value of $1/4$, then as long as $q > .04$, he would reject the two simple joint propositions just cited and conclude prior to experimentation that e_H & $h_{.99} \vee e_T$ & $h_{.01}$.

Even an objectivist necessitarian can reach this conclusion provided he has background knowledge of the sort I have described.

But once X has reached this conclusion, he may argue that he should simulate the behavior prescribed by B_2 . However,

he need not argue that he should use a program for employing his or Y 's observation report as input. If he waits until he can add his report or Y 's to his corpus before choosing a gamble, the report added will entail the truth of either $h_{.99}$ or $h_{.01}$ and will uniquely decide which of the options is E -admissible for him. There will be no need to invoke considerations of S -admissibility.

Reasoning like this will work neatly enough only in special cases. It will not be compelling in the original cases 4 or 5 or more complicated variants of these.

However, something very much like the case where the alternatives are $h_{.99}$ and $h_{.01}$ obtains in situations where X is designing a program he anticipates using repeatedly in a large number of similar situations. For example, X might anticipate having to choose between gambles like G_1 and G_2 in a large sequence of contexts where he is presented with a coin which either has a .6 or a .4 chance of landing heads and where he has an opportunity to observe a toss of the coin before accepting the gamble (as in case 5).

Suppose that X is constrained to follow a policy of using the same program over and over again and that he is concerned to reap the greatest net benefit from repeated application in the long run.

X does not know what kind of coin will be presented to him on each occasion. It could be either a coin with a .6 or a .4 chance of landing heads. There are, therefore, 2^n distinct possibly true hypotheses concerning the sequence of such coins he will face—given that the sequence has a length n . There are $n + 1$ distinct hypotheses as to the relative frequency of .6 coins in the sequence.

Given any specific hypothesis as to the relative frequency of .6 coins, X can compute a chance distribution for the relative frequency of heads among the tosses of .6 coins and for the relative frequency of heads among the tosses of .4 coins. By direct inference, he can assign, if he is an objectivist necessitarian, an interval-valued credence to the proposition obtained by conjoining the hypothesis as to the relative frequency of .6 coins with a specification of relative frequency of heads on tosses of the .6 coins and of relative frequency of heads on tosses of the .4 coins. Let U consist of all such

hypotheses and let the M -function assign equal value to all elements of U .

Even for low values of q , X will be justified in rejecting elements of U which assert (a) that r/n coins in the sequence are .6 coins, (b) that the percentage of heads in the r cases where a .6 coin is tossed deviates from 60% by a large amount, and (c) the percentage of heads in the $n - r$ cases where a .4 coin is tossed deviates from 40% by a large amount.

By bookkeeping through reiterating the rejection procedure, the conclusions can be strengthened. (See section 2.7.) Once such expansion has been made, then if X has only the four programs B_1 , B_2 , B_3 , and B_4 for repeated use, the average long-run benefits of using one of these programs when $n = r$ will be approximately equal to the expectation of that program conditional on $h_{.6}$. When $n = 0$, it will be approximately equal to the expectation conditional on $h_{.4}$. The average benefits for intermediate values of r will fall between these values.

If X is interested in long-run benefits and fixes security levels according to method (a), repeatedly applying B_1 , B_2 , and B_4 would all be E -admissible and P -admissible, but only B_2 would be S -admissible. There would be no issue as to whether an alternative method of fixing security levels is to be used.

The alternative policies have to be construed as programs for repeated use of observation reports as inputs. If the reports are used as evidence, the methods described could not support repeated use of B_2 (i.e., simulating conformity to B_2).

Thus, there is a long-run rationale for adopting B_2 that avoids relying on Wald's criterion or variants of it for each single case. It may seem, therefore, that we have found a rationale which avoids converting moral prescriptions into dictates of reason while rendering objectivist necessitarianism free from difficulty.

That is not so. The rationale works only when X seeks to promote net long-run benefits—or equivalently average long-run benefits—rather than with the benefits involved on each occasion of use.

Consequently, objectivist necessitarianism cannot be defended by appealing to long-run considerations unless it is claimed that men would be irrational unless they were concerned with long-run interests and not with the short run. Appeal to the long run has been made by a great many phi-

losophers; but I balk at the notion that someone is irrational if he adopts a more myopic view.

In section 17.1, the benefits of routine expansion and decision making were noted. But even routinization was seen as a benefit not because the long-run benefits are being promoted but because it minimizes various costs on each application and, in the case of routine expansion, enhances our capacities for using data in acquiring new information.

Since routine programs are used where, so we assume, the error probabilities are very low, we could anticipate prior to each use of the routine that the observation report made will be true. Hence, if upon making a report, the report made conflicted with what is in the corpus, we would not escape contradiction by refusing to add the report to the corpus.

In practice it will make little difference whether we expand the corpus via induction prior to each application of the routine or not. The need to obtain new information via the testimony of the senses plus the benefits of routinization will suffice to warrant adopting a routine provided the error probability is low enough.

But it is neither feasible nor desirable to routinize all decision making. We are, indeed, creatures of custom and habit most of the time. But we design experiments precisely in those cases where custom and habit have failed us.

Proponents of the Neyman-Pearson-Wald school of thought attempt to convert all decision making to programmed decision making where the agent is designed to respond in appropriate ways to observation reports. Wherever possible, the program is designed to be a routine to be applied over and over again in the hope of realizing some long-run benefit. Such ideas are not peculiar to Neyman, Pearson, and Wald. They are to be found in the writings of C. S. Peirce, among others.

I contend that we often cannot and when we can sometimes should not obey the prescriptions of such an approach to decision making. In particular, in scientific inquiry, we are concerned to obtain new error-free information pertaining to the particular problem or cluster of problems under investigation. It is small consolation to be told, for example, that Neyman-Pearson methods of confidence interval estimation

will, in the long run of applications, only rarely lead to error. In inquiry, as in daily life, we are interested in the errors which might obtain in the specific context under scrutiny.†

† The interpretation of confidence interval estimation as implementing a routine program for using observation reports as inputs appears to capture the views of Neyman and Pearson. Some authors (e.g., A. Birnbaum, "Concepts of Statistical Evidence," in *Philosophy, Science, and Method: Essays in Honor of Ernest Nagel*, edited by S. Morgenbesser, P. Suppes, and M. White, New York: St. Martin's, 1969, pp. 121-124 and D. R. Cox and D. V. Hinkley in *Theoretical Statistics*, London: Chapman and Hall, 1974, pp. 45-46 and 48-49) construe the calculation of a confidence interval as a summary of the evidential relevance of the data. They favor this conception of evidential import over others (e.g., likelihood) because it can be operationalized with reference to the long-run reliability of hypotheses favored by the evidence according to this concept. Nothing said in this chapter applies to these views. I do not understand them except in a formal way. For me a concept of evidential support becomes intelligible only if its function in inquiry and deliberation can be identified. This sort of "operationalization" seems to me to be crucial. Long-run sampling properties have only an obscure relevance to concepts of evidential support.

Frege opposed psychologism in deductive logic. He contended that laws of logic are not "psychological laws of takings-to-be-true but laws of truth."¹ Furthermore, "what is true is something objective and independent of the judging subject; for psychological logicians it is not."²

Frege did concede, however, that although logic does not describe how agents think, it does prescribe how they ought to think:

It will be granted by all at the outset that the laws of logic ought to be guiding principles for thought in the attainment of truth, yet this is only too easily forgotten, and here what is fatal is the double meaning of the word "law." In one sense a law asserts what is; in the other it prescribes what ought to be. Only in the latter sense can the laws of logic be called "laws of thought"; so far as they stipulate the way in which one ought to think. Any law asserting what is, can be conceived as prescribing that one ought to think in conformity with it, and is thus in that sense a law of thought. This holds for laws of geometry and physics no less than for laws of logic. The latter have a special title to the name "laws of thought" only if we mean to assert that they are the most general laws, which prescribe universally the way in which one ought to think if one is to think at all.³

Frege apparently supposed that once we are given the laws of logic construed as objective truths, we may readily derive prescriptions regulating how agents ought to think. Consequently, there is little point in directing attention to prescriptions concerning the way one ought to think additional to that focused on the study of logical truth. This view is reiterated by Carnap, who applied it to inductive logic as well as deductive logic. An objectivist view of logic construes principles of logic as objective truths "independent of the contingency of the facts of nature" and "not dependent upon whether or what any person may happen to imagine, think, believe, or know about these sentences."⁴ A qualified psychologistic view of

logic sees logic as regulating the way agents ought to think. According to Carnap, "psychologism thus diluted has virtually lost its content; the word 'thinking' or 'believing' is still there, but its use seems gratuitous."⁵

Frege and Carnap could not have intended to claim that prescriptions about the way agents ought to think are deducible from logical truths. The logical truths must be supplemented by some principle or system of principles licensing the derivation of such prescriptions from logical truths. Given these principles, the study of objective logical truth would suffice for the determination of norms of rational thought. Unfortunately neither Frege nor Carnap are clear concerning what these principles are like.

Consider, for example, whether *X*'s body of beliefs at *t* should be deductively closed. If the beliefs are consciously held beliefs, no one can meet the requirement even approximately. For this reason, the prescription may be restricted in application to ideally rational agents. Alternatively one might follow the approach I favor and impose the deductive closure requirement not on an agent's consciously held beliefs but on propositions the agent is committed to accepting as evidence in deliberation and inquiry. Thus, the prescription may be imposed on agents who lack perfect memory, computational facility, and emotional health. But even this suggestion is fraught with controversy. Kyburg denies that we should be committed to deductively closed sets of hypotheses. To be sure, Kyburg's notion of acceptance or belief does not appear to be acceptance as evidence. But the existence of controversy over the domain of applicability of a requirement of deductive closure points precisely to what is wrong with the remarks of Frege and Carnap on what Carnap calls qualified psychologistic construals of logic. It is not a trivial matter to specify the overarching norms which link the objective truths of logic with prescriptions about how agents ought to think. Neither Frege nor Carnap are explicit about the domain of applicability of the laws of thought normatively construed. Nor do they indicate what form such laws ought to take.

Nonetheless, there is a kernel of truth in what they say about qualified psychologism. The principles of qualified psychologistic deductive logic should be derivable from the truths of objective deductive logic and general principles "which

prescribe universally the way in which one ought to think if one is to think at all.”

These general principles should be applicable no matter who the agents are or the circumstances under which they are deliberating; and, given these general principles, the truths of objective logic should determine the normative laws of thought. Thus, for both Frege and Carnap, qualified psychologistic logic imposes constraints on the way agents ought to think which are universally applicable regardless of circumstance.

I have followed Frege and Carnap in regarding normative principles regulating deliberation and inquiry to be principles of logic only if the norms are universally applicable and relatively context independent. This view of qualified psychologistic logic has been employed both in relation to deductive and to inductive logic. I do not object to this aspect of the Frege-Carnap attitude towards qualified psychologism but to the casual manner in which they address the problem of deriving normative laws of thought from objectivist logic.

What remains to be considered is whether there are other context-dependent norms regulating inquiry and deliberation additional to those context-independent principles of logic. The only norms Frege and Carnap seem to countenance are context-independent principles. Once context is taken into account, critical control seems to vanish. Contextual factors may be investigated by psychology, sociology, and history; but how such factors regulate deliberation and inquiry is beyond critical control. Consequently, we must say that either some aspect of scientific inquiry or of other forms of deliberation is subject to objective critical control in the sense that it is regulated by principles of qualified psychologistic logic or it is not subject to critical control at all.

The pervasive influence of this polarity is exhibited in the tendency of students of rational probability judgment to divide into two camps. Necessitarians seek to defend the view that all rational agents should be obliged to make the same probability judgments relative to the same evidence. In chapters 14–17, several efforts to save necessitarianism have been considered and found wanting.

The typical response to the deficiencies of necessitarianism is to adopt some form of personalism which prescribes that

rational agents should choose numerically precise credal states subject only to the relatively weak constraints of credal coherence and, perhaps, direct inference.

Popper shared Carnap's concern to construe scientific inquiry and deliberation as regulated by objective principles of criticism in the sense opposed to psychologism.⁶ Unlike Carnap, however, Popper was an extreme subjectivist concerning probability judgment. On Popper's view, it seems that even credal coherence may be violated.⁷ Precisely because probability judgment is beyond critical control, it cannot play a central role in scientific inquiry. Otherwise science would be infected with subjectivity. The objectivity of science is protected by decontaminating it of the poisons of Bayesianism. Within Popper's theory, only the principles of deductive logic regulate thought.

But, even within his own circle, Popper's outlook has been subjected to trenchant criticism. P. K. Feyerabend has argued that any methodological constraints of a fixed and context-independent variety which may be acceptable are too weak to impose significant restrictions on deliberation and inquiry. Just as intemperate personalists insist that anything goes (except for violation of credal coherence, consistency, and uniqueness) in the choice of credal states, Feyerabend claims that anything goes in our efforts to undermine and refute one theory through confrontation of that theory with others.⁸

Thus, both within the Bayesian tradition exemplified by Carnap's work on probability and within the Popperian tradition, the objectivity of scientific inquiry is made to stand or fall with the existence of a fairly powerful and fixed system of principles applicable to all agents on all occasions. If such a system exists, the objectivity of scientific inquiry and knowledge is saved. If it does not, we must see scientific inquiry as buffeted by psychological, sociological, and historical factors in a manner beyond serious critical control.

This polarization is not a new one; but it is interesting that both Carnap and Popper have seen their commitments to objectivity as deriving from Frege's attack on psychologism, and that much contemporary discussion concerning scientific method proceeds on assumptions about the terms of debate advanced in their writings even when the particular outlooks of Carnap and Popper are rejected.

For this reason, I regard anyone to be suffering from the curse of Frege who submits to the polarization and chooses either in favor of method and against psychologism, sociology, and historicism or chooses against method and in favor of psychologism, sociology, and historicism. Frege himself is not particularly responsible for spreading the curse in the form I have described; but his views on psychologism in deductive logic do seem to have been elaborated upon or modified in ways central to the polarization. If the curse were, counter to fact, a blessing, few would object to honoring Frege by calling it after his name. There is little room for complaint in naming the curse in his honor.

Not everyone suffers from the curse. A. Shimony, for example, has sought to identify relevant factors of the context of inquiry such as the problem under investigation and the potential answers to the question under serious scrutiny in devising critical principles for the determination of credal states additional to the weak advice which inductive logic can offer. Within the Popperian tradition, I. Lakatos sought to indicate how features of a research program might direct the conduct of specific inquiries whose results might in turn lead to modification and eventual abandonment of the program itself. It is unnecessary to endorse the details of either Shimony's or Lakatos's proposals in order to register approval of the general thrust of their approach. As pragmatists have recognized for a long time, contextual considerations may be invoked in appraising steps taken at various stages of inquiry. We need not rest content with deductive logic and the thin gruel that inductive logic has to offer.

Exhortation alone will not exorcise a curse. One must seek to construct an alternative to suffering from its tyranny. In *Gambling with Truth*, I began constructing such an account by making some proposals concerning inductive expansion in certain special cases. In this book, I have sought to extend this account of inductive expansion and supplement it with an account of routine expansion and of contraction so that, at least in outline, a view of how revisions of bodies of knowledge or evidence construed as standards for serious possibility should be evaluated has been constructed.

In the later chapters of the book, I have tried to integrate this view of the revision of knowledge with an account of the

revision of probability judgment which, like the proposals concerning standards for serious possibility, accords the contexts of specific inquiries an important place in critical appraisal of changes in cognitive commitment.

Even if the proposals made here could be judged adequate without further modification or elaboration, more needs to be done. Factors like potential answers, informational value, problems worth investigating, and the like have been recognized to be relevant to the appraisal of revisions of cognitive commitments. I have not, however, explored the extent to which they are subject to critical control in inquiry and deliberation. The topic of abduction has not as yet been adequately addressed.

But even if a suitable account of abduction can be constructed congenial with the outlook I have been advocating, it is doubtful whether context alone can guarantee definite conclusions and definite probability judgments with the level of precision we often desire. We should be prepared to acknowledge that we ought often to suspend judgment pending further inquiry rather than arbitrarily leap to conclusions. We should also acknowledge that inquiries occur in history, so that attention to context is important and that the knowledge we use and the methods we employ in such inquiries are themselves subject to revision in the course of these inquiries.

Perhaps there is some incorrigible knowledge. Perhaps there are a few fixed methodological principles constraining the revision of knowledge and probability judgment. I have considered some fixed features of method in this book and, if the conclusions reached here are sound, these fixed principles are very weak.

This circumstance ought not lead, however, to skepticism about the objectivity of science or of its methods. What passes for scientific method at the moment depends on the current corpus of scientific knowledge and the current methods for appraising hypotheses with respect to probability. Our methods are modified by our knowledge just as our knowledge is modified according to our methods.

In revising our knowledge and methods, we are not enthralled by frameworks, discourses, or other conceptual hobgoblins from which we can escape only by shifting to other apparitions of a similar character and only in a manner which

precludes reasoned appraisal of the rival frameworks from a vantage point which begs no questions. We are not prevented from seeking error-free information in a rational way relative to the best information available; and we can open our minds to rivals to current doctrine without succumbing to conversion.

It is, therefore, possible to keep revisions of our changing doctrines and methods under critical control provided we are prepared to recognize the relevance of contextual factors to such control and to recognize that the relevance of context (and, in this sense, of historical circumstance) to the appraisal of progress does not threaten us with an objectionable psychologism, sociologism, or historicism. There is, indeed, very little fixed method; but, with all due respect to those who suffer from the curse of Frege, there is objectivity enough.

APPENDIX A BRIEF SERMON ON ASSESSING ACCIDENT RISKS IN U.S. COMMERCIAL NUCLEAR POWER PLANTS

Advocates of the use of nuclear energy have urged us to compare the expected gains and losses deriving from the use of nuclear power with the expected gains and losses resulting from alternative approaches to energy supply. The public is encouraged to take into account how remote the probability of a serious accident is and consider whether the expected benefits adequately compensate for the risks incurred. Invoking the Bayesian principle to maximize expected utility, proponents of nuclear energy dismiss the anxieties of those who worry about the worst possible consequences of extensive reliance on energy generated in nuclear power plants as based on irrational or foolish appeals to the principle of minimizing the worst possible consequences (or maximizing the minimum possible benefit).

But maximin is not always foolish. Calculations of expected utility require numerical evaluation of society's utilities and identification of numerically precise estimates of the probabilities of accidents of various sorts. If the available evidence fails to warrant a sufficiently definite system of credal probability judgments for use in computing expected utilities, considerations of expected utility may fail to render a verdict concerning the merits of the alternatives of permitting and promoting nuclear power, prohibiting it, or more refined variant approaches to our energy problems. When neither permitting nor prohibiting the extensive use of nuclear power is ruled out by appeal to considerations of expected utility, it is perfectly reasonable to look at the worst possible consequences of each available option which has survived the test of expected utility (each *E*-admissible option) and identify that option or those options for which the worst possible consequence (the security level) is better (or at least no worse) than the worst possible consequence of all other *E*-admissible op-

tions. Maximizing comes into its own when maximizing fails to render a verdict.

By and large the public seems to endorse values which recognize the worst possible consequences (the security level) of permitting a nuclear power plant to run to be worse than the worst possible consequences of prohibiting such a plant. Perhaps this assumption is mistaken or unclear. And perhaps we should not listen to society. These are difficult and vexing issues. I shall not attempt to settle them here.

In any case, critics of extensive use of nuclear power do think the security level of a nuclear policy to be far inferior to the security level of prohibiting its use. Advocates of nuclear power do not question that assumption. Instead they belittle its relevance. We beg no questions by taking it for granted here.

If the security level for refusing to run a nuclear plant is greater than the security level for permitting its use, the only condition under which permitting a nuclear power plant seems cogent is when the probability of serious accident is sufficiently low according to all permissible expectation-determining probability distributions (i.e., all distributions not justifiably eliminated on the basis of the available evidence) to warrant the verdict that the option of running the plant is optimal with respect to expected utility according to all permissible distributions (i.e., is uniquely *E*-admissible).

If that could be established, the circumstance that prohibiting the plant bears a higher security level would not count against the policy of running the plant.

Thus, advocates of nuclear power are right to appeal to calculations of expected utility as relevant to the assessment of rival policies. On the other hand, to make the case that such calculations are decisive, advocates must argue that no matter what expectation-determining or credal probability distribution permitted by the evidence is used, running nuclear plants bears greater expected utility than prohibiting the use of such plants.

Critics of nuclear power have a somewhat easier job. They do not have to show that prohibition bears greater expected utility than promotion. Success in establishing this would be sufficient to make their case but not necessary. If prohibiting the running of nuclear power plants bears maximum expected

utility according to some expectation-determining probability distributions permitted by the evidence and permitting the running of such plants bears maximum expected utility according to other such distributions, considerations of expected utility fail to decide and maximin may be invoked. According to our assumptions about society's values, prohibiting the nuclear plants will be recommended.

Much debate concerning energy policy and the role of nuclear power plants does appear to polarize around advocates of nuclear power who argue like expected utility maximizers and critics who argue like maximinners. The approach to rational choice advocated in this book and outlined in the previous paragraphs denies that this dispute reveals a disagreement over public values or over principles of rational choice. Perhaps such differences are involved but this is not evident. The heart of the disagreement concerns assessments of the probabilities of accidents of a serious nature occurring at nuclear power plants. Are the probabilities of accidents of a serious nature sufficiently low according to all permissible distributions to favor nuclear power or is the verdict sufficiently indeterminate to justify prohibiting extensive use of nuclear power because of considerations of security?

In 1975, the Nuclear Regulatory Commission published a report entitled *Reactor Safety Study* which sought to give an assessment of probabilities of serious accidents in commercial nuclear reactors. In the first two appendices to this study, the methods employed in making the evaluations of probabilities are described.¹ There is considerable discussion of the difficulties involved in making such assessments and of the methods employed to meet the difficulties.

One important task is to identify the various types of episodes which could lead to a major accident. Given the identification of such episodes, the possible ways in which the backup and safety devices which are installed to prevent such episodes from leading to serious consequences need to be identified, and the probabilities that these devices will function as designed (or alternatively the probabilities that they will fail) need to be estimated.

Unfortunately the available data on the reliability of backup and safety devices provide a poor basis for making the requisite estimates of chances. Information may be available con-

cerning the performance of some of these devices at various plants; but the conditions of operation vary widely from plant to plant, so that one cannot assume that the chance of breakdown of a system in a given interval of time will be the same for one plant as will the chance of breakdown of a corresponding system in another plant.

The backup systems and safety devices (such as sources of alternative electrical supply, alternative cooling systems, etc.) are complex systems. The data base for deriving estimates of the chance of such a system breaking down might, therefore, be broadened by examining data for estimating failure probabilities for critical components of the system and calculating a failure probability distribution for the entire system from this information.

To do this, one must be able to represent the condition of failure of the entire system as a Boolean function of the failure conditions of the critical components. In appendix 2 of *Reactor Safety Study*, the method of fault tree analysis is introduced as a method of graphic representation of system failure as a Boolean function of failures of components.

Thus, in an example used for illustrative purposes in appendix 2 (called program SAMPLE), seven distinct components are identified such that failure of the total system depends on appropriate distributions of failures among these components.

If T represents the event of system failure and $X(n)$ represents failure of the n th component (where $1 \leq n \leq 7$), program SAMPLE represents T as

$$(1) T \text{ iff } [X(1) \vee X(6) \vee X(7) \vee [(X(2) \vee X(3)) \& (X(4) \vee X(5))]].$$

The term "fault tree" derives from the use of tree diagrams as graphical representations of Boolean functions.

There are 2^7 possible ways in which the seven components can be in states of failure and nonfailure. For some of these T holds (i.e., the system has failed) and for some it does not. The probability that T holds is the sum of the probabilities of the members of that subset of the 2^7 possible states of the seven components for which T holds.

In general, empirical data are lacking concerning the probabilities of these 2^7 states (or the comparable 2^n states of other systems of components for other complex systems). However,

some data may be available concerning the failure probabilities of each of the seven components separately. The authors of *Reactor Safety Study* described methods for simplifying computations so that calculations of system failure probabilities could be made from failure probabilities of the seven components. For the most part, the simplifications tended to increase failure probabilities and, hence, load the case against the safety of the reactors under investigation. For example, the probability of $X(1) \vee X(6)$ was taken to be the sum of the probabilities for $X(1)$ and $X(6)$. The probability of $X(1) \& X(6)$ was not deducted from this, even though in a strict probabilistic calculation it should be. On the other hand, in the example, the authors of the study did assume that $X(3)$ and $X(5)$ are probabilistically dependent (for illustrative purposes) and used a simplification which tended to reduce the probability of failure. Despite this, I shall employ the method of simplification adopted in *Reactor Safety Study* for determining system probability of failure from failure probabilities of components. The simplified formula is

$$(2) P(T) = (P(X(2) + P(X(3)))(P(X(4)) + P(X(5))) + P(X(1)) + P(X(6)) + P(X(7)) + P(X_{cm})).$$

Here X_{cm} is the joint failure of the third and fifth component.

Setting aside worries one might have about the methods of simplification involved (which in my view are marginal worries), it should be pointed out that the analysis could overlook some potential source of breakdown of the system under study. The authors of the report are aware of this and attempt to make allowance for it and for human failure. In any case, failure to account for all serious and relevant possibilities threatens all deliberate decision making. This does not mean that we should not attempt to be as clear as we can concerning such possibilities. It means only that we should avoid congratulating ourselves too much and looking at our analyses as security blankets. I do not think the authors of *Reactor Safety Study* are to be faulted on this score.

The failure rate for a given component at a given time t is the probability of that component breaking down in the interval from t to $t + dt$ for extremely small dt (or, more accurately, an approximation to that probability).

The authors of the study assumed that critical components

would be subject to regular test and maintenance. Under this assumption, the failure rate for a given component should be approximately constant over time and independent of the past background of breakdowns. In other words, the behavior of the component with respect to failure should exemplify a Poisson process. The probability density $f(t)$ characterizing the probability distribution for t being the time of first breakdown starting from initial time 0 should approximate $\lambda e^{-\lambda t}$, and the constant failure rate should be λ .

If failure rates for all components could be ascertained with reasonable precision, we might utilize (2) to obtain a failure rate for the entire system.

But even if a failure rate can be imputed on the basis of the available data to the operation of a given component in the system operating in a specific plant, the assessment of the failure rate for the corresponding component of a corresponding backup system in another plant may and, indeed, is, in general, revealed to be different; for the conditions of operation in different plants have an impact on failure rates (which, though they may not be small enough, are still small enough to be sensitive to variations in operating conditions relevant to calculations of expected losses due to accidents.)

Thus, the available data show a spectrum of failure rates for a given type of component in a given type of system. If we seek an expected failure rate for a similar component in a new system of the same type or an average failure rate for all such systems in a planned network of nuclear power plants, it is desirable to ascertain the objective or statistical or chance probability distribution over values of the failure rate of the type of component falling in the detected or "assessed" interval of values for that failure rate. The new system (or the systems belonging to the planned network of nuclear power plants) is treated as if it were selected at random from a population of systems. The systems already in operation from which data are collected are also treated in the same manner and an effort is made on the basis of that data to ascertain the chance distribution of failure rates for the type of component in the population sampled.

Notice that if a definite objective statistical or chance distribution could be established on the basis of such data at least to a good degree of approximation, direct inference could then

license a uniquely permissible credal probability function to be used in computing expected losses from serious accidents should a nuclear plant be set in operation or should an entire network of such plants be set in operation.

Unfortunately the range of failure rates assessed for a component on the basis of the data often range over two orders of magnitude—e.g., from .01 to .0001. If differences in failure rates of a given component of one in a thousand might matter to the safety of a given component or backup system, data even from a fairly large number of components of a given type might prove insufficient to exploit standard chi-square goodness-of-fit tests effectively. Furthermore, the best that such tests can do is provide some approximate determination as to whether some hypothesis about the chance distribution of failure rates ought to be ruled out on the basis of the data or is in adequate agreement with the data. If the data are too sparse, it may turn out that many diverse distributions "fit" the data.

The authors of *Reactor Safety Study* acknowledge the sparsity of the data and that such data are insufficient to favor a definite chance distribution. Nonetheless, the authors do adopt a definite distribution over the failure rates for each component investigated. They are prepared to construe that distribution as an objective chance or frequency distribution in some contexts and as a "Bayesian" or credal distribution in others. I take them to mean that they are prepared to reach a conclusion as to what statistical probability distribution should be used as the basis for a direct inference to determine a credal distribution for use in computing expectations.

The authors of *Reactor Safety Study* begin their search for a distribution for the failure rate λ by deriving an "assessed range" of values for λ from the data. They then determine that log normal distribution for λ which entails a 5% probability of λ taking a value below the lower bound of the assessed range and a 5% probability of its taking a value greater than the upper bound of the assessed range.

For example, if the assessed range of values of λ is from .00125 to .02, the log normal distribution fitted to this interval according to the method just described has a mode of .00245, median of .005, and mean value of .00715.

These measures of location are nearer the lower end of the

assessed range of failure rates than to the upper end, as is typical of the log normal distribution.

Let ν be the sum of the lower and upper bounds of the assessed range and $\kappa = \nu - \lambda$. The value of κ ranges over the same interval of values as the value of λ . It takes its maximum value when λ takes its minimum value and vice versa. The authors of *Reactor Safety Study* could easily have fitted a log normal distribution for κ . In our example, $\nu = .02125$ and the mode, median, and mean value for κ would be equal to the corresponding values for λ when λ is assumed log normal. But now the distribution of λ is such that the mode is .0188, the median is .01625, and the mean value is .0141.

Thus, adopting a log normal distribution for κ tends to concentrate more probability in the upper end of the assessed range for λ than in the lower range, in sharp contrast to adopting a log normal distribution for λ .

If data were abundant, we might look forward to evidence supporting the elimination of one or the other of these distributions as not adequately in agreement with the data. The authors of *Reactor Safety Study* do not consider how well the log normal for κ fits data, and I have not explored their data tables to make a determination. One suspects, however, from their own testimony that in many cases the data were too sparse to warrant ruling out either of these distributions even though the measures of location differ substantially in one case from the other.

The authors of the report do appear to be mindful of the thin ice on which they are skating. They keep emphasizing that the data do not conflict with their use of the log normal distribution and that in those few cases where there were enough data for a chi-square test, the log normal for λ is "consistent" with the data. But, of course, the same data could be consistent with the log normal for κ as well. I am not especially fond of the log normal for κ . Any other distribution diverging from the log normal distribution for λ could do just as well.

The authors note that

The log normal distribution is thus a "natural" distribution for describing data which can vary by factors in the same way that a normal distribution is "natural" when the data can vary by additive or subtractive increments. In the study, one of the

reasons that the log normal was assessed to be suitable was that the component and other input data in general could vary by factors. For example, a failure rate estimated at 10^{-6} could vary from 10^{-7} ($= 10^{-6}/10$) to 10^{-5} ($= 10^{-6} \cdot 10$).²

I am not sure why data which can vary by factors cannot also vary by increments. Most assuredly they do. I assume that the authors are alluding to a variety of formal properties of log normal distributions which render them attractive for purposes of computation of the sort they subsequently explain in somewhat greater detail.

In any case, these formal properties (or nearly all of them) belong to the log normal distribution for κ as well as to the log normal distribution for λ and do not explain the bias in favor of the latter distribution. There is one consideration invoked which does not apply to both distributions equally well:

The log normal distribution form, in particular its positive skewness, can incorporate general reliability-associated behaviors of the assessed data (the positive skewness accounts for the occurrence of less likely but large deviate values, such as abnormally high failure rates due to batch defects, environmental degradation, and other outlier causing effects).³

Please observe that this argument on behalf of the log normal for λ presupposes that the failure rates at the upper end of the assessed range are abnormal. The authors do not actually say that the data establish this. They simply assume it. Given this assumption, skewing the distribution in the opposite direction as the log normal for κ does is ruled out of consideration at the outset.

The authors of *Nuclear Safety Study* comment on their arguments for the log normal for λ as follows:

The above items of course do not constitute tenets for the dogmatic justification of the log normal as the only distribution applicable, but instead serve as a priori considerations. As a complement to the above considerations, from a pragmatic point of view the log normal was employed because it was flexible, it was consistent with reliability and data properties and it is a standardly employed and straightforward (null hypothesis) distribution. Checks and tests of its applicability to the data of this study do not contradict nor refute these a priori and pragmatic considerations.⁴

Finally the authors defend the care and objectivity with which they chose the log normal by comparing it with another

distribution of failure rates—the log uniform over the assessed range for λ . They found “insignificant differences” over the assessed ranges because of their size in using these two distributions and, hence, regarded the choice of the log normal as “robust.” We note, of course, that the log uniform is even more skewed in favor of low failure rates than the log normal.

In spite of the bizarre character of some of the arguments advanced by the authors of *Reactor Safety Study*, I am not prepared to suggest that authors have chosen the log normal distribution precisely because it will load the dice in favor of a positive verdict concerning the reliability of backup systems for nuclear power plants. And I detect no basis in what I have read for the accusation that their judgments are politically or economically motivated. They present their reasons, no matter how distressingly poor they might be, with a frankness which suggests that the authors of the report regard the reasons as perfectly good ones.

What is troubling about the procedures described by these authors is that the methods they employ for assessing probabilities of failure conform to a wide variety of epistemological outlooks favored by contemporary philosophers and students of probabilistic inference. I think a great many epistemologists would be hard pressed to explain why, from their point of view, there is anything wrong with the choice of the log normal for λ .

I do not see anything objectionable in pointing out that the data do not rule out that distribution. But there is something seriously wrong in choosing that distribution to guide conduct when there is no warrant for doing so.

According to personalist Bayesians, we cannot expect too much help from the data in determining a uniquely permissible probability distribution in guiding our conduct. We are, nonetheless, under an obligation to adopt a credal state satisfying credal uniqueness—i.e., according to which one probability distribution is uniquely permissible. As long as the distribution favored by the authors of the report satisfies coherence requirements, there is nothing to complain about in the arbitrary choice of that distribution. What is needless, perhaps, is the elaborate effort to defend that choice by bad argument where no argument is needed.

The authors of the report, however, did seem anxious to

ground judgments of credal probability on judgments as to which of rival statistical hypotheses concerning the objective chance distribution of failure rates is correct. In the face of insufficient data, the authors did not conclude that they should suspend judgment between the rival statistical hypotheses and look to their credal state for the various seriously possible rivals for help in determining (via direct inference) the credal state to adopt. If the view I have developed in the present book is correct, that approach would have led to an indeterminate state of credal judgement concerning failure rates.

Instead, the authors of *Reactor Safety Study* reasoned that when more than one rival hypothesis “fits” the data reasonably well, one is free to choose between the rivals by an appeal to “a priori” (which I interpret in this context to mean “formal” and “aesthetic”) reasons or “pragmatic” considerations.

They are not alone in finding such reasoning acceptable. It is Quine, after all, who observes that in fitting curves to data we should pick the simplest of the many curves agreeing with the data, acknowledging all the while that simplicity is to a considerable degree a question of taste. I see little difference between such appeals to simplicity and the appeals to “naturalness” in *Reactor Safety Study*. I detect little difference between the appeals to customary practice by the authors and appeals to the “taste for the familiar” or invocations of the tradition of some branch of inquiry in its “normal” phase. (The authors of *Nuclear Safety Study* appeal to what they apparently think is a paradigmatic application of the log normal.)

One should not blame the personalists, Quine, or anyone else in the large list of distinguished authors whose epistemologies do not prohibit the analysis carried out in *Nuclear Safety Study* for the excesses of that study. But it is surely worthwhile to reevaluate epistemological outlooks incapable of rendering a negative verdict on the practice of evaluating risks of a serious accident in nuclear plants in the manner of the authors of the study under review.

The moral of the story is that we should learn to suspend judgment. We should, in the case under consideration, learn to acknowledge that the data justifies and, indeed, obligates us to suspend judgment concerning the objective chance distribution over failure rates within a given range of values. We

should be prepared to adopt credal states for hypotheses about failure rates in specific cases which are indeterminate and which allow many diverse distributions to be permissible.

This does not mean that we should become total skeptics. Scientific inquiry has furnished us with much knowledge and, in some contexts, with information which justifies appraising risks and expectations with a considerable degree of precision. But although we should prize precision when we can get it, we should never pretend to precision we lack; and we should be ever mindful of our ignorance even when it hurts.

I do not know whether a reconsideration of the data contained in *Reactor Safety Study* in keeping with respect for our ignorance as well as for what we know would vindicate those who seek to halt our reliance on nuclear sources of energy or whether the proponents of nuclear energy would be shown to be in the right.

However, it may be helpful to compare the analyses of program SAMPLE reported by the authors of *Reactor Safety Study* with an analysis of their illustrative data in accordance with principles I find relevant to the problem at hand.

Median failure rates are given (for illustrative purposes) for the seven factors in program SAMPLE together with an "error factor" which determine the upper and lower bounds of the assessed range of failure rates for each factor.

One method used to obtain a system failure rate was to substitute the median values for the failure rates for the $X(n)$'s (on the assumption of log normal distributions) and the corresponding median value for X_{cm} (reckoned as $\sqrt{a^2b}$, where a is the median for $X(3)$ and b is the median for $X(5)$) into formula (2). This led to an estimated system failure rate of .0056. (The range of values for $X(7)$ is from .00001 to .0000001 and offers negligible contribution to the system failure rate. It was therefore omitted from computations in formula (2).)

This estimate was then contrasted with a more sophisticated appraisal of system failure rates. A program for selecting values in the assessed range of a given input factor at random (or simulating such selection) on the assumption that values are distributed in accordance with the log normal distribution for that factor. This is done for each of the input factors. The values are substituted into (2) and a system value is obtained. The process is repeated a great many times (1,200 times) and

a range of system of values is derived. A log normal distribution is then fitted to this range according to the method described previously and the median, upper, and lower values identified.

The results for program SAMPLE give an upper bound of .0144, a median of .0078 and a lower bound of .0045. The median value obtained is thus higher than the result obtained by the first method.

Keep in mind, however, that these results are obtained from the arbitrary choice of log normal distributions of failure rates for the several inputs. In the absence of any warrant for picking out any coherent probability distribution over failure rates in the assessed range for a given component, all such distributions should be permissible. If the result is a family of probabilities of serious accident in the reactor which fails to decide between permitting and prohibiting operation of the reactor on grounds of expected utility maximization, maximin should be invoked.

This means, however, that a relevant assessment of the failure rate of the system under investigation (in our example it is program SAMPLE) need take into account only the system failure rate determined by substituting the upper bound values for the failure rates of the inputs. In program SAMPLE, one should substitute into formula (2) the largest value of the failure rate for $X(n)$. For X_{cm} , this should be $\max(X(3), X(5))$.

The system failure rate computed on this basis is .0564.

The result is nearly four times as great as the upper bound of the range of system failure rates obtained by the authors of *Reactor Safety Study* using Monte Carlo methods of propagation (the second method described above) and nearly eight times as great as the median value determined in that way.

If failure rates for containment and backup systems were appraised in the manner I propose, we would not have to pretend to knowledge we do not have. And if probabilities of serious accidents remained sufficiently low, calculating failure rates in the manner indicated, the arguments of proponents of nuclear energy that the reactors are safe enough would be more compelling.

Of course, better data could lead to sharper verdicts concerning failure rates. And improved technology might lead to

reconsideration of pessimistic assessments of the safety of nuclear plants.

In the wake of recent episodes at nuclear power plants, not only have policies favoring nuclear energy been discredited in the eyes of the public; but the credibility gap has extended to pure scientists and engineers. I do not seek to join the chorus of critics of scientific approaches to the solution of problems. If the serious problems we face are to be solved, they will be solved only by careful and sophisticated inquiry.

But the demands of such inquiry require that experts admit the limits to what they know. Scientists and technologists should not pretend to a knowledge they do not have because a government or a public demands that they be supplied with answers to questions for which there is insufficient evidence. And the public and government should understand and respect the limits on what they can expect of responsible scientists and engineers. They should refrain from putting unreasonable pressures on investigators to subvert their better judgment.

Above all we should beware of epistemologies which permit us to violate this counsel and indulge our tastes for the familiar, simplicity, explanatory power, naturalness, paradigmatic methods of puzzle solving, and the like without regard to the risks of error both in theory and in practice which our indulgences may be incurring.

REFERENCES

Chapter 1

- 1 C. S. Peirce, *The Philosophy of Peirce*, edited by J. Buchler. New York: Harcourt Brace, 1950, p. 58.
- 2 *Ibid.*, p. 59.
- 3 *Ibid.*, pp. 46-47.
- 4 *Loc. cit.*
- 5 *Op. cit.*, pp. 38-39. This is a passage from the essay "How to Make Our Ideas Clear."
- 6 K. R. Popper, *Conjectures and Refutations*, New York: Basic Books, 1962, pp. 223-237.
- 7 W. V. O. Quine, *Word and Object*, New York: Wiley and MIT, 1960; Cambridge, Mass.: MIT Press, 1964, p. 25.
- 8 Popper, *op. cit.*, pp. 238-240.

Chapter 2

- 1 C. W. Peirce, *Collected Papers*, Cambridge, Mass.: Harvard University Press, 1931, v. 1, p. 154.
- 2 W. V. O. Quine and J. Ullian, *The Web of Belief*, New York: Random House, 1970, p. 43.
- 3 *Ibid.*, p. 55.
- 4 W. V. O. Quine, *The Ways of Paradox*. New York: Random House, 1966, pp. 242-245. See also Quine and Ullian, *op. cit.*, p. 46.
- 5 The ideas outlined in this and the remaining sections of this chapter were initially introduced in my *Gambling with Truth*, New York: Knopf, 1967; Cambridge, Mass.: MIT Press, 1973. My proposals in that essay were modified and summarized in "Information and Inference," *Synthese*, v. 17 (1967), pp. 369-391. For an excellent critical review of my proposals and related proposals of others, see R. Hilpinen, *Rules of Acceptance and Inductive Logic*, Amsterdam: North-Holland, 1968.
- 6 Philosophers unfamiliar with measure theory in general and Lebesgue measure in particular may find the summary offered by R. C. Jeffrey in *Studies in Inductive Logic and Probability* (edited by R. Carnap and R. C. Jeffrey. Berkeley: University of California Press, 1971, v. 1, pp. 169-221) helpful. H. E. Kyburg covers similar material in a useful manner in an appendix to his *Logical Foundations of Statistical Inference* (Dordrecht: Reidel, 1974).

7 The notion of a density function is briefly explained by both R. C. Jeffrey and H. E. Kyburg in the works cited in ref. 6.

8 "Information and Inference," pp. 369-391.

9 I. Levi, "Acceptance Revisited," in *Local Induction*, Dordrecht: Reidel, 1976, pp. 1-71.

10 *Gambling with Truth*, pp. 84-85.

11 *Ibid.*, pp. 149-152, adumbrates this point.

12 J. Hintikka, "A Two Dimensional Continuum of Inductive Methods," in *Aspects of Inductive Logic*, edited by Hintikka and P. Suppes, Amsterdam: North-Holland, 1966, pp. 113-132; H. Jeffreys, *Theory of Probability*, 3rd ed., London: Oxford University Press, 1961, pp. 128-132; and A. Shimony, "Scientific Inference," in *The Nature and Function of Scientific Theories*, edited by R. G. Colodny, Pittsburgh: University of Pittsburgh Press, 1970.

Chapter 3

1 E. Whittaker, *A History of the Theories of Aether and Electricity*, New York: Harper, 1960, v. 1, pp. 386-387, 390-391.

2 T. Kuhn, "Reflections on my Critics," in *Criticism and the Growth of Knowledge*, edited by I. Lakatos and A. Musgrave, Cambridge University Press, 1970, p. 262.

3 *Ibid.*, pp. 265-266.

4 P. K. Feyerabend, "Consolations for the Specialist," in *Criticism and the Growth of Knowledge*, pp. 197-230, and *Against Method*, London: NLB, 1975.

5 *Ibid.*, p. 265.

Chapter 4

1 B. De Finetti, "Foresight: Its Logical Laws, Its Subjective Sources," translated by H. E. Kyburg in *Studies in Subjective Probability*, edited by H. E. Kyburg and H. E. Smokler, New York: Wiley, 1964, p. 111.

2 L. J. Savage, *The Foundations of Statistics*, New York: Wiley, 1954, p. 61. Savage does not flatly preclude the introduction of stronger principles; but his skepticism is apparent.

3 H. Jeffreys, *Theory of Probability*, 3rd ed., London: Oxford University Press, 1961, and *Scientific Inference*, 2nd ed., Cambridge: Cambridge University Press, 1957.

4 R. Carnap, *Logical Foundations of Probability* 2nd ed., Chicago: University of Chicago Press, 1962, and *Continuum of Inductive Methods*, Chicago: University of Chicago Press, 1952.

5 De Finetti, *op. cit.*

6 Savage, *op. cit.*

7 R. Carnap, "The Aim of Inductive Logic," in *Logic Methodology and Philosophy of Science*, edited by E. Nagel, P. Suppes, and A. Tarski, Stanford: Stanford University Press, 1962, pp. 305-316; "In-

ductive Logic and Rational Decisions" and "A Basic System of Inductive Logic, Part I," in *Studies in Inductive Logic*, edited by R. Carnap and R. C. Jeffrey, Berkeley: University of California Press, 1971, pp. 1-166.

8 J. M. Keynes, *A Treatise on Probability*, London: MacMillan, 1921, pp. 33-34.

9 B. O. Koopman, "The Bases of Probability," *Bull. American Math. Soc.*, v. 46 (1940), pp. 763-774.

10 I. J. Good, "Subjective Probability as the Measure of a Nonmeasurable Set," in *Logic Methodology and Philosophy of Science*, pp. 319-329.

11 C. A. B. Smith, "Consistency in Statistical Inference and Decision" (with discussion), *J. Royal Stat. Soc.*, ser. B, v. 23 (1961), pp. 1-25.

12 F. Schick, *Explication and Inductive Logic*, doctoral dissertation, Columbia University, 1958.

13 H. E. Kyburg, *Probability and the Logic of Rational Belief*, Middletown, Conn.: Wesleyan University Press, 1961, and *The Logical Foundations of Statistical Inference*, Dordrecht: Reidel, 1974.

14 A. R. Dempster, "Upper and Lower Probabilities Induced by a Multivalued Mapping," *Annals Math. Stat.*, v. 38 (1967), pp. 325-339.

15 H. Jeffreys, *Theory of Probability*, p. 37, and R. Carnap, *The Logical Foundations of Probability*, pp. 219-226.

16 R. A. Fisher, *Statistical Methods and Statistical Inference*, 2nd edition, New York: Hafner, 1959, pp. 31-35.

17 E. S. Pearson, "Some Thoughts on Statistical Inference" *Annals Math. Stat.*, v. 33 (1962), pp. 394-403; reprinted in *The Selected Papers of E. S. Pearson*, Berkeley: University of California Press, 1966. See in particular pp. 277-278.

18 H. Reichenbach, *Experience and Prediction*, Chicago: University of Chicago Press, 1938, pp. 3-6, 317.

19 H. E. Kyburg, *Logical Foundations of Statistical Inference*, ch. 10.

20 L. J. Savage, *The Foundations of Statistics*, p. 61.

21 A. Shimony, "Scientific Inference," *Pittsburgh Studies in Philosophy of Science*, v. 4, edited by R. G. Colodny, Pittsburgh: University of Pittsburgh Press, 1970, pp. 79-172, especially sec. III.

22 I. J. Good, "Rational Decisions," *J. Royal Stat. Soc.* ser. B, v. 14 (1952), p. 114.

23 See the works cited in ref. 7.

Chapter 5

1 R. C. Jeffrey, *The Logic of Decision*, New York: McGraw Hill, 1965, pp. 8-9 and R. Nozick, *The Normative Theory of Rational*

Choice, unpublished doctoral dissertation, Princeton University, 1963.

2 F. P. Ramsey, "Truth and Probability," in *The Foundations of Mathematics and Other Essays*, New York: Humanities Press, 1950, reprinted in *Studies in Subjective Probability*, edited by H. E. Kyburg and H. Smokler, New York: Wiley, 1964, pp. 69-82 and B. De Finetti, "Foresight: Its Logical Laws, Its Subjective Sources," translated by H. E. Kyburg in *Studies in Subjective Probability*, pp. 100-110.

3 A. Shimony, "Coherence and the Axioms of Confirmation," *J. Symbolic Logic*, v. 20 (1955), pp. 8-20.

4 R. Carnap, "The Aim of Inductive Logic," *Logic Methodology and Philosophy of Science*, edited by E. Nagel, P. Suppes, and A. Tarski, Stanford: Stanford University Press, 1962, pp. 305-316.

5 B. De Finetti, "On the Axiomatization of Probability Theory," in *Probability, Induction and Statistics*, New York: Wiley, 1972, p. 91 and footnote added by De Finetti for the translation.

6 Ramsey, *op. cit.*, p. 80.

7 De Finetti, "Foresight," pp. 109-110.

8 Shimony, *op. cit.*, p. 7.

9 De Finetti, "Axiomatization," pp. 81-82.

10 I. Levi, "Coherence, Regularity and Conditional Probability," *Theory and Decision*, v. 9 (1978), pp. 10-11. The rule QRECU is called QCEU in this article.

11 *Ibid.*

12 De Finetti, "Axiomatization," p. 99.

13 I rely on P. Halmos, *Measure Theory*, New York: Van Nostrand (1950) for the concepts and results of measure theory used in this discussion.

14 De Finetti, "Axiomatization," pp. 90-94.

15 H. Jeffreys, *Theory of Probability*, 3rd ed., London: Oxford University Press, 1961, pp. 115-125.

16 I. Hacking, *Logic of Statistical Inference*, Cambridge: Cambridge University Press, 1965, p. 153.

17 Halmos, *op. cit.*, p. 31.

18 A. Renyi, *Foundations of Probability*, San Francisco: Holden Day, 1970, ch. 2 and especially p. 40.

19 This example is taken from De Finetti, "Axiomatization," pp. 99-100.

20 D. V. Lindley, *Bayesian Statistics, A Review*, Philadelphia Society for Industrial and Applied Mathematics, 1971, pp. 50-51; A. P. Dawid, M. Stone, and J. V. Zidek, "Marginalization Paradoxes in Bayesian and Structural Inference," *J. Royal Stat. Soc.*, ser. B, v. 35 (1973), pp. 189-223, with discussion; and M. Stone, "Strong In-

consistency from Uniform Priors," *J. American Stat. Assoc.*, v. 71 (1976), pp. 114-116, followed by responses from several distinguished statisticians including Lindley. The significance of Lindley's recantation will be appreciated by those familiar with his *Introduction to Probability and Statistics*, Cambridge: Cambridge University Press, 1965, part 2.

21 See De Finetti, "Axiomatization," pp. 98-99.

Chapter 6

1 I. Levi, *Gambling with Truth*, New York: Knopf, 1967; Cambridge, Mass.: MIT press, 1973. See especially the first 10 chapters.

2 I. Levi, "Information and Inference," *Synthese* v. 17 (1967), pp. 369-391. See, in particular, n. 34 for a statement of a major revision of the account of epistemic utility found in *Gambling with Truth*. Note should also be taken of n. 33.

3 I. Levi, "Acceptance Revisited," in *Local Induction*, edited by R. J. Bohdan, Dordrecht: Reidel, 1976, pp. 1-71.

4 *Gambling with Truth*, pp. 47-49.

5 *Ibid.*, pp. 82-85. Further discussion of the rule for ties is found in my "Truth, Content and Ties," *J. Phil.*, v. 68 (1971), pp. 865-876.

6 *Gambling with Truth*, *loc. cit.*

7 This idea was first suggested by me in "Potential Surprise in the Context of Inquiry", in *Uncertainty and Expectation in Economics, Essays in Honor of G. L. S. Shackle*, edited by C. F. Carter and J. L. Ford, Oxford: Blackwell, 1972, p. 230. See also "Acceptance Revisited," p. 56.

Chapter 7

1 R. D. Luce and H. Raiffa, *Games and Decisions*, New York: Wiley, 1958, p. 13.

2 G. L. S. Shackle, *Expectation in Economics*, Cambridge: Cambridge University Press, 1949, ch. 2.

3 C. A. B. Smith, "Consistency in Statistical Inference and Decision," *J. Royal Stat. Soc.*, ser. B., v. 23 (1961), pp. 1-25.

4 A. Wald, *On the Principles of Statistical Inference*, Notre Dame, Ind.: University of Notre Dame Press, 1942, ch. 6 and especially pp. 44-45.

5 J. Von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*, Princeton: Princeton University Press, 1944.

Chapter 8

1 For a forceful statement of this point of view, see M. Friedman and L. J. Savage, "The Utility Analysis of Choices Involving Risks," *Landmarks in Political Economy*, edited by E. J. Hamilton, A. Rees, and H. G. Johnson, Chicago: University of Chicago Press, 1962, pp. 326-327.

2 This approach is developed at great length for utility theory in P. C. Fishburn, *Decision and Value Theory*, New York: Wiley, 1964.

pp. 6-14, 60-71, and 83-88. I. J. Good endorses a similar approach to credence in "Rational Decisions," *J. Royal Stat. Soc.*, ser. B, v. 14 (1952) and "Subjective Probability as the Measure of a Non-measurable Set," in *Logic Methodology and Philosophy of Science*, edited by E. Nagel, P. Suppes, and A. Tarski, Stanford: Stanford University Press, 1962, pp. 319-329.

3 Good adopts a similar approach for credal probabilities in "Subjective Probability as the Measure of a Non-measurable Set," p. 327. I am considering the parallel approach for utility here.

4 This is Good's view in *loc. cit.*

5 Savage disparages views such as this in *The Foundations of Statistics*, New York: Wiley, 1954, p. 58.

6 Good and Fishburn are concerned with these matters.

7 This, at any rate, reflects the approach of H. E. Kyburg, "Conjunctivitis," *Induction, Acceptance and Rational Belief*, edited by M. Swain, Dordrecht: Reidel, 1970, pp. 55-82.

8 See Fishburn, *op. cit.*, pp. 83-88. Fishburn denies that consequences can be incomparable unless incomparability is analyzed in terms of uncertainty, dishonesty, limitations of computational facility, or the like. Fishburn's view is vivid testimony to Bayesian insensitivity to value conflict.

9 See J. Harsanyi, *Essays on Ethics, Social Behavior, and Scientific Explanation*, Dordrecht: Reidel, 1976, pp. 10-13.

10 I. Levi, *Gambling with Truth*, New York: Knopf, 1967; Cambridge, Mass.: MIT Press, 1973, pp. 149-152.

11 I. Levi, "Abduction and Demands for Information," *The Logic and Epistemology of Scientific Change*, edited by I. Niiniluoto and R. Tuomela, *Acta Philosophica Fennica*, v. 30, nos. 2-4, North-Holland, 1979, pp. 405-429.

Chapter 9

1 See G. L. S. Shackle, *Expectation in Economics*, Cambridge: Cambridge University Press, 1949; 2nd ed., 1952; *Decision, Order and Time*, Cambridge: Cambridge University Press, 1961; 2nd enlarged ed., 1969. For the reconstruction of Shackle's view of potential surprise alluded to here, see I. Levi, "On Potential Surprise," *Ratio*, v. 8 (1966), pp. 107-129; *Gambling with Truth*, New York: Knopf, 1967; Cambridge, Mass.: MIT Press, 1973, chs. 8 and 9; "Potential Surprise in the Context of Inquiry," in *Uncertainty and Expectation in Economics: Essays in Honor of G. L. S. Shackle*, edited by C. F. Carter and J. L. Ford, Oxford: Blackwell, 1972, pp. 213-236.

2 See *Gambling with Truth*, ch. 9. L. J. Cohen has recently introduced a measure of inductive probability having the formal properties of what I called a measure of degrees of confidence of acceptance, $b(g)$, where $b(g) = d(\sim g)$. Moreover, Cohen has recognized the relevance of such a measure to Keynes' notion of weight of argument.

See L. J. Cohen, *The Probable and the Provable*, Oxford: Clarendon Press, 1977.

3 B. O. Koopman, *The Bases of Probability*, BAMS, v. 46 (1940), pp. 763-774; reprinted in H. E. Kyburg and H. E. Smokler, *Studies in Subjective Probability*, New York: Wiley, 1964, pp. 161-172.

4 I. J. Good, "Subjective Probability as the Measure of a Non-measurable Set," in *Logic Methodology and Philosophy of Science*, edited by E. Nagel, P. Suppes, and A. Tarski, Stanford: Stanford University Press, 1962, pp. 319-329.

5 C. A. B. Smith, "Consistency in Statistical Inference and Decision," *J. Royal Stat. Soc.*, ser. B, v. 23 (1961), pp. 1-37.

6 H. E. Kyburg, *Probability and the Logic of Rational Belief*, Wesleyan University Press, 1961.

7 F. Schick, *Explication and Inductive Logic*, doctoral dissertation, Columbia University, 1958, pp. 63-77.

8 D. H. Krantz, R. D. Luce, P. Suppes, and A. Tversky, *Foundations of Measurement*, New York: Academic Press, 1971, pp. 200-201.

9 P. Suppes, "The Measurement of Belief," *J. Royal Stat. Soc.*, ser. B, v. 36 (1974), pp. 160-191, with discussion.

10 *Ibid.*, pp. 188-189.

11 *Ibid.*, pp. 168-171.

12 *Ibid.*, p. 169.

13 Koopman, *op. cit.*

14 S. Spielman, "Bayesian Inference with Indeterminate Probabilities," *PSA 1976*, v. 1 (1976), pp. 185-196. See p. 189, theorem 4.

15 *Ibid.*, pp. 192-194.

16 R. D. Luce and H. Raiffa, *Games and Decisions*, New York: Wiley, 1957, pp. 288-290.

Chapter 10

1 H. E. Kyburg does just this in "Randomness and the Right Reference Class," *J. Phil.*, v. 74 (1977), pp. 519-520.

2 W. L. Harper, "Rational Belief Change, Popper Functions, and Counterfactuals," *Synthese*, v. 30 (1975), pp. 221-262.

Chapter 11

1 I. Levi and S. Morgenbesser, "Belief and Disposition," *American Phil. Quarterly*, v. 1 (1964), pp. 221-232. This essay is reprinted in a useful anthology edited by R. Tuomela (*Dispositions*, Dordrecht: Reidel, 1978).

2 The crucial passages from Carnap's "Testability and Meaning" are reprinted in Tuomela, *op. cit.*, pp. 3-16.

3 This point is made, albeit in a form I would not use, by D. H. Mellor in "In Defense of Dispositions," reprinted in Tuomela, *op. cit.*, pp. 71-72.

4 D. K. Lewis, "Probabilities of Conditionals and Conditional Probabilities," *Phil. Review*, v. 85 (1976), pp. 297-315.

Chapter 12

1 R. Von Mises, *Probability, Statistics, and Truth*, London: Allen and Unwin, 1957, p. 30.

2 B. De Finetti, "Foresight: Its Logical Laws, Its Subjective Sources," *Studies in Subjective Probability*, edited by H. E. Kyburg and H. Smokler, New York: Wiley, 1964, pp. 113-118.

3 I. Levi, *Gambling with Truth*, New York: Knopf, 1967; Cambridge, Mass.: MIT Press, 1973, pp. 219-221.

4 B. De Finetti, *Probability, Induction and Statistics*, New York: Wiley, 1972, p. 99.

5 M. Stone, "Strong Inconsistency from Uniform Priors," *J. American Stat. Soc.*, v. 71 (1976), pp. 114-116.

6 De Finetti, *Probability, Induction and Statistics*, pp. 98-100. The article cited therein was originally published in Italian in 1949.

7 Stone, *op. cit.*

8 Most notably D. V. Lindley. See his response to Stone's paper in the discussion that follows the paper.

9 De Finetti, *Probability, Induction and Statistics*, p. 104.

Chapter 13

1 R. A. Fisher, *Statistical Methods and Scientific Inference*, 2nd ed., New York: Hafner, 1959, pp. 19-20.

2 *Ibid.*, pp. 31-34.

3 J. Neyman, "Outline of a Theory of Statistical Estimation Based on the Classical Theory of Probability," *Phil. Trans. Royal Soc. London*, ser. A., v. 236, no. 767 (1937), pp. 333-380; reprinted in *A Selection of Early Papers of J. Neyman*, Berkeley: University of California Press, 1967, pp. 250-290. See p. 259.

4 E. S. Pearson, "Some Thoughts on Statistical Inference," *Annals Math. Stat.*, v. 33 (1962), pp. 394-403; reprinted in *Selected Papers of E. S. Pearson*, Berkeley: University of California Press, 1966, pp. 276-283. See pp. 278-279.

5 R. A. Fisher, "The Arrangement of Field Experiments," *J. Ministry of Agriculture of Great Britain*, v. 33 (1925), pp. 504-505.

6 *Ibid.*, pp. 506-507.

7 S. Spielman, "Levi on Personalism and Revisionism," *J. Phil.*, v. 72 (1975), pp. 785-803.

8 W. Edwards, H. Lindeman, and L. J. Savage, "Bayesian Statistical Inference for Psychological Research," *Psychological Review*, v. 70 (1963), pp. 201-208.

9 I. Levi, "On Potential Surprise," *Ratio*, v. 8 (1966), p. 115; I. Hacking, *Logic of Statistical Inference*, Cambridge: Cambridge University Press, 1965, pp. 89-91. I proposed the initial likelihood

rejection rule. Hacking proposed a likelihood rejection rule which is derivable from mine by reiteration.

10 A. Shimony, "Scientific Inference," *Pittsburgh Studies in Philosophy of Science*, v. 4, edited by R. G. Colodny, Pittsburgh: University of Pittsburgh Press, 1970, pp. 97-103.

Chapter 14

1 R. A. Fisher, "Inverse Inference," *Proc. Cambridge Phil. Soc.*, v. 26 (1930), pp. 528-535 and especially p. 533.

2 R. A. Fisher, *Statistical Methods and Scientific Inference*, 2nd ed., New York: Hafner, 1959, p. 51.

3 J. Neyman, "On the two different aspects of the representative method: the method of stratified sampling and the method of purposive selection," *J. Royal Stat. Soc.*, v. 97 (1934), pp. 558-625; reprinted in *A Selection of Early Papers of J. Neyman*, Berkeley: University of California Press, 1966, pp. 98-141, including discussion. The remarks of R. A. Fisher are of special interest. In subsequent publications, Neyman gave up attempting to tame fiducial inference. (Neyman, "Fiducial Argument and the Theory of Confidence Intervals," *Biometrika*, v. 32 (1941), pp. 128-50; reprinted in *Selection*, pp. 375-394.)

4 H. Jeffreys, *Theory of Probability*, 3rd ed., London: Oxford University Press (1961), p. 381.

5 I. Hacking, *Logic of Statistical Inference*, Cambridge University Press, 1965, pp. 138-148.

6 J. W. Tukey, "Some Examples with Fiducial Relevance," *Annals Math. Stat.*, v. 28 (1957), pp. 687-695.

7 T. Seidenfeld, *Philosophical Problems of Statistical Inference: Learning from R. A. Fisher*, Dordrecht: Reidel, 1979, chs. 4 and 5.

Chapter 15

1 R. A. Fisher, *Statistical Methods and Scientific Inference*, 2nd ed., New York: Hafner, 1959, p. 74.

2 I. Levi, *Gambling with Truth*, New York: Knopf, 1967; Cambridge, Mass.: MIT Press, 1973, p. 101.

3 I. Hacking, *Logic of Statistical Inference*, Cambridge: Cambridge University Press, 1965, pp. 32-38.

4 *Ibid.*, p. 135.

5 *Ibid.*, ch. V.

6 *Ibid.*, pp. 141-143 and 149-150.

7 *Ibid.*, p. 62.

8 *Ibid.*, p. 135.

9 *Ibid.*, pp. 141-143.

10 *Ibid.*, pp. 66-71.

11 *Ibid.*, pp. 149-150.

12 *Ibid.*, pp. 146 and 147-148.

13 *Ibid.*, pp. 152-154.

14 *Ibid.*, p. 148.

15 T. Seidenfeld, *Philosophical Problems of Statistical Inference: Learning from R. A. Fisher*, Dordrecht: Reidel, 1979, chs. 4 and 5.

16 D. V. Lindley, "Fiducial Distributions and Bayes' Theorem," *J. Royal Stat. Soc.*, ser. B, v. 20 (1958), pp. 102-107.

17 This construal of fiducial distributions is suggested by the remarks of D. A. S. Fraser in "On Fiducial Inference," *Annals Math. Stat.*, v. 32 (1961), pp. 668-669.

Chapter 16

1 R. A. Fisher, "Mathematical Probability in the Natural Sciences," *Technometrics*, v. 1 (1959) pp. 21-29; reprinted in v. 5 of *Collected Papers of R. A. Fisher*, edited by J. H. Bennet, Adelaide: University of Adelaide Press, 1974, pp. 398-406.

2 R. A. Fisher, *Statistical Methods and Scientific Inference*, 2nd ed., New York: Hafner, 1959, pp. 32-33, 51-57, and 110-111.

3 Such as "The Nature of Probability," *Centennial Review*, v. 2 (1958), pp. 261-274, reprinted in *Collected Papers*, v. 5, pp. 384-397; and "Scientific Thought and the Refinement of Human Reasoning," *J. Operations Research Soc. Japan*, v. 3 (1960), pp. 1-10, reprinted in *Collected Papers*, v. 5, pp. 475-484.

4 Fisher, "Mathematical Probability in the Natural Sciences," pp. 23-24.

5 *Ibid.*, pp. 21-23.

6 R. B. Braithwaite, *Scientific Explanation*, Cambridge: Cambridge University Press, 1953, pp. 126-129.

7 Fisher, "Mathematical Probability in the Natural Sciences," pp. 22-23.

8 *Ibid.*, p. 23.

9 *Loc. cit.*

10 *Ibid.*, pp. 23-24.

11 H. Reichenbach, *Experience and Prediction*, Chicago: University of Chicago Press, 1938, pp. 316-319; *The Theory of Probability*, Berkeley: University of California Press, 2nd ed., 1971, pp. 372-378.

12 H. E. Kyburg, *The Logical Foundations of Statistical Inference*, Dordrecht: Reidel, 1974, pp. 286-290.

13 "New Methods for Reasoning Towards Posterior Distributions Based on Sample Data," *Annals Math. Stat.*, v. 35 (1966), pp. 364-372; and "A Generalization of Bayesian Inference," *J. Royal Stat. Soc.*, ser. B, v. 30 (1968), pp. 208-219.

14 H. E. Kyburg, "Randomness and the Right Reference Class," *J. Phil.*, v. 74 (1977), pp. 501-521.

Chapter 17

1 J. Neyman and E. S. Pearson, "On the Problem of the Most Efficient Tests of Statistical Hypotheses," reprinted in *Joint Statistical Papers of J. Neyman and E. S. Pearson*, Berkeley: University of California Press, 1966, pp. 141-142. This paper originally appeared in *Phil. Transactions Royal Soc.*, ser. A, v. 231 (1933), pp. 289-337.

2 I. Hacking, *Logic of Statistical Inference*, Cambridge: Cambridge University Press, 1965, p. 104.

3 E. S. Pearson, "Some Thoughts on Statistical Inference," reprinted in *Selected Papers of E. S. Pearson*, Berkeley: University of California Press, 1966, p. 277. This paper originally appeared in *Annals Math. Stat.*, v. 33 (1962), pp. 394-403.

4 Neyman and Pearson, *op. cit.*, p. 142.

5 A. Wald, *On the Principles of Statistical Inference*, Notre Dame, Ind.: University of Notre Dame Press, 1942, in particular ch. 6.

Chapter 18

1 G. Frege, *The Basic Laws of Arithmetic: Exposition of the System*, a translation of part of *Grundgesetze der Arithmetik* by M. Furth, Berkeley: University of California Press, 1967, p. 13.

2 *Ibid.*, p. 15.

3 *Ibid.*, p. 12.

4 R. Carnap, *Logical Foundations of Probability*, 2nd ed., Chicago: University of Chicago Press, 1962, p. 38.

5 *Ibid.*, p. 41.

6 K. R. Popper, *Objective Knowledge*, Oxford: Clarendon Press, 1972, p. 162.

7 *Ibid.*, pp. 78-79.

8 P. K. Feyerabend, *Against Method*, London: Verso, 1975, p. 23 and throughout the book.

Appendix

1 *Reactor Safety Study. An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants: Appendix 1. Accident Definition and Use of Event Trees* (108 pages) and *Appendix 2. Fault Trees*, Washington, D.C.: Nuclear Regulatory Commission, Oct. 1975.

2 *Op. cit.*, *Appendix 2*, p. 42.

3 *Ibid.*, p. 42.

4 *Ibid.*, p. 43.

INDEX

- Abduction, 42-46, 49-50
 Ability, 236-237
 adverbially modified, 241-243
 and compulsion, 239
 and serious possibility, 244-246
 Acceptance, 132
 Admissible option, 96
 Analytic reasoning. *See* Corrigibility
 Anomaly, 62-63
 Awareness, 10
- Bar-Hillel, Y., 48n
 Bayes, T., 297, 315
 Bayesian conditions or doctrine;
 Bayesianism, 51, 78, 84, 87, 94,
 95-96, 98-100, 132, 133, 135,
 136, 150, 152, 164-168, 207-210,
 229
 and conflict, 168-210
 Bayes' theorem, 83-85, 125, 129,
 236, 289, 325
 Belief, justified, 1
 Birnbaum, A., 423n
 Bolzano, B., 287
 Borel, E., 287
 Braithwaite, R. B., xv, 234n, 370
- Cantor, G., 287
 Carnap, R., 48n, 67, 78, 79, 80n, 86,
 89, 99, 238, 243, 297, 375, 424-
 427
 Caution, 50-56, 137, 181
 throwing caution to the winds, 56
 CD-predicate, 248-249
 Chance, 48, 86-87, 234-236, 250-287
 and credal coherence, 261
 and credal consistency, 261-264
 and credal probability, 234
 existence of, 264-271, 392-394
 Kyburg's view, 376
 and the long run, 271-274
 predicate, 250-252
 composite, 253
 simple, 252
 semantics of, 256-261
 frequency interpretation, 256-261,
 394
 set-up, 37n, 236
 on single trials, 271-273
 and repeatable trials, 273-274
 subjunctive conditionals supported
 by, 276-278
- CIL, 86, 89-90
 Classical representation of a
 decision problem, 105-106
 Cognitive decision problems, 51
 Cognitive resource, 74ff.
 Coherentism, 86
 Commitment, 164-168, 186-188
 Compatibility of evaluations of
 expected utility, 116
 strong condition of, 114
 Compulsion, 237
 Conditionalization
 confirmational, 81, 83-86, 216-217,
 222-225, 230-233, 289, 297-298,
 383-385
 and contraction, 82n
 Dempsterian (*D*-
 conditionalization), 385-392
 inverse temporal credal, 81
 of one credal state by another, 80
 temporal credal, 81, 289
 Conditionals, 246-248
 Confidence interval estimation, 326,
 422, 423n
 Confirmational commitment, 48, 79-
 83, 216-217, 289
 and credibility, 80n
 logical, 86
 Confirmational irrelevance, 225-233
 Confirmational tenacity, 82, 86, 87,
 99, 216

Conglomerability, 118, 280
 finite, 278
 Context dependence, 47, 56, 65-67,
 91-93, 315, 322, 426-430
 Contraction, 25, 33, 58, 59-62
 Contradiction, 62-63
 Corpus of knowledge, 1-2, 7, 74ff.
 contradictory, 27-28
 hierarchy of corpora, 9
 potential, 12-13
 Corrigibility, 1-5, 13-19, 59, 99
 categorical, 7-8
 conceptual, 7-8
 degrees of, 61-62
 idiosyncratic, 7
 Cox, D. R., 423n
 Credal
 coherence, 77-78, 104, 118-119,
 144, 261
 consistency, 78, 104, 118-119, 261-
 264
 convexity, 78-79, 191-214
 and the multiplication theorem,
 193-196
 intervalism, 197-204, 212, 375-376
 Dempsterian, 387-392
 irrelevance, 107
 regularity, 77, 110-116
 states, 74
 uniqueness, 79, 84, 87
C(UK), 82
 Cumulative distribution function,
 327
 Curse of Frege, 318-319, 424-430

 Decision making
 deliberate, 35-37
 under risk, 98
 routine, 35-37
 theoretical and practical, 72-76
 under uncertainty, 98, 144-147
 De Finetti, B., 78, 86, 99, 109, 110,
 112, 114, 115, 116, 118, 122, 126,
 129-130, 131, 210, 219n, 285,
 287, 392-394
 Demands for information, 34, 46-47,
 180-182
 Dempster, A. P., 79, 102, 200, 297-
 298, 325, 386n, 387n, 388n, 369,
 385-392
 Density function, 327
 Dewey, J., xvi

 Direct inference, 86, 144, 236, 250-
 287
 from composite hypotheses
 finite case, 278-280
 countably infinite case, 280-284
 continuous case, 284
 Fisher's view, 369-375
 Kyburg's principle of, 374n, 375-
 385
 and pedigree theories of
 knowledge, 396-398
 principle of, 254-256, 405n
 Reichenbach's view, 373-374
 Disposition, 237-239
 semantics for, 243-244
 Dominance
 strong, 106
 and credal coherence, 108-109
 principle of (*SD*), 107
 weak, 106, 109-112
 principle of (*WD*), 109

E-admissibility
 of options, 96-97, 144-145
 principle of, 96, 104
 Edwards, A. W. F., 344n
 Epistemic utility, 50-51, 132, 180-
 182
E-undominated option, 136
 Evidential import, 423n
 Expansion, 25, 33, 34-35
 deliberate or inferential, 34-40,
 137, 428-429
 and statistical prediction, 274-276
 and likelihood rules of rejection,
 311-315
 routine, 34-40, 99
 via appeals to witnesses, 37, 51-
 57, 400-403
 via observation, 37, 399-423
 Expected utility, 95-96
 conditional, 112-113
 epistemic, 135-136
 maximizing, 51-53, 132
 principle of, 96
 Experimental densities, 356-358

 Fallibilism, 8, 13-19, 34-35, 41, 57
 categorical, 14-15, 18
 epistemological, 13-14, 18-19, 70-
 72
 and free speech, 31-32

Feyerabend, P. K., 34, 58, 68-70,
 427
 Fiducial inference, 102, 325, 326,
 340n
 commitment to irrelevance, 328,
 331-334, 338-339
 inversion step, 328, 330-331, 337
 Kyburg's theory, 382-383
 pivotal step, 328, 329-330, 336-337
 principle of
 continuous case, 339
 countably infinite case, 335-336
 finite case, 334
 and insufficient reason, 335, 336,
 340
 as summarizing data, 367-368
 tame, 326, 342
 Fisher, R. A., 87, 90, 102-103, 294,
 296, 301-302, 316, 325, 335, 341,
 342, 343, 368-376, 382, 383
 Focal pair, 148
 Frege, G., 424-428
 Frequency principle, 346, 405n

 Gettier problem, 29-30
 Good, I. J., xvi, 79, 96, 186, 187,
 199, 204, 208, 375

 Hacking, I., 37n, 102, 126, 234,
 242n, 297, 326, 327, 329, 345-
 349, 354, 356-358, 360-363, 368,
 404-405
 Harper, W., 219
 Harper functions, 219-221
 Harsanyi, J., 149n, 150n, 175
 Higgenbotham, J., xvi
 Hilpinen, R., 48n
 Hinkley, D. V., 423n
 Hintikka, J., 48n, 56
 Hope level, 148
 Hurwicz, L., 145n

 Ignorance, 183-186, 315
 Bayesian, 184
 cognitive, 183
 strong, 184
 credal, 185
 and cognitive ignorance, 185
 Incommensurability, 67-68
 Independence of irrelevant
 alternatives, 208-210
 Inductive logic, 77-78, 85-88, 98

 Informational value, 34, 45-48, 50-
 51
 and contraction, 59-60
 Insufficient reason, 184, 315
 Irrelevance. *See K*-irrelevance, *L*-
 irrelevance, Confirmational
 irrelevance, Credal irrelevance,
 Stochastic irrelevance
 principle of, 346, 347, 354-355

 Jeffrey, R. C., 52, 82n, 106
 Jeffreys, H., 56, 78, 79, 89, 90, 99,
 126, 127, 129, 130, 297, 326, 360,
 375
 Joint proposition, 347-348

 Keynes, J. M., 79, 90, 183, 375
K-irrelevance, 374n, 377, 382-383,
 395-396
 Kolmogorov, A. N., 110, 346, 346n
 Koopman, B. O., xv, 79, 199, 206-
 207, 346, 375
 Knowledge
 pedigree theory of, 1-2
 as a standard for serious
 possibility, 2-7, 71-72, 74
 as true justified belief, 1-3, 28-30
 Krantz, D. H., 204
 Kuhn, T., 34, 58, 65-70, 71
 Kyburg, H. E., xv, 79, 102, 199,
 203, 297-298, 325, 369, 374n,
 375-385, 386n, 392-398

 Lakatos, I., 428
 Laplace, P. S., 315
 Lehman, E. L., 409-410
 Lehrer, K., 54n
 Leximin, 145
 admissibility, 149, 211
 and spreads in the odds, 154-156
 optimality, 149
 Likelihood, 84, 228, 236, 289, 343-
 368, 423n
 assumption of uniqueness, 229
 continuous case, 355-358
 and deriving posteriors from
 priors, 343-344
 discrete cases, 347-352
 and fiducial inference, 352-354,
 358-359
 its function in inquiry, 345, 367-
 368
 HL-likelihood, 348

Likelihood (*cont.*)
 law of, 102
 and insufficient reason, 359-363
 as a principle of inductive logic
 grounded in our conception of
 chance, 361-363
 refuted, 366-367
 and rejection, 311-315
 Lindley, D. V., 130, 340n, 363, 366
L-irrelevance, 377, 382-383, 395-396
 Long run, 419-423
 Lorentz, H. A., 63
 Luce, R. D., 145, 204, 209, 228n

Maximin, 98, 145. *See also* Security
 as pessimistic, 149-151
 Maxwell, J. C., 62
 Measuring credence, 204-208
M-function. *See* Probability,
 information-determining
 Michelson, A. A., 62-63
 Mill, J. S., 32
 Minimax, 134. *See also* Maximin
 Minimax regret, 145
 Mixed option, 162-163, 410-413
 Morgenbesser, S., xvi, 234-235n,
 237-238
 Morgenstern, O., 163
 Multiplication theorem extended to
 the continuous case, 123-125

Nagel, E., xvi
 Necessitarianism, 89, 90-91, 93, 426
 Neyman, J., 37-38n, 87, 90, 102,
 156-157, 293-296, 326, 404-409,
 415, 419, 422-423
 Normal science, 67-69
 Nozick, R., 106
N-predicate, 239
 explicit, 239-240

Objectivist inductive logic, 87, 91,
 102, 144, 279-280
 and determination of priors and
 likelihoods, 289-292
 Objectivist necessitarianism, 102-
 103, 280, 292-296
 and the relevance of data, 296-298,
 403-408
 and routine decision making, 421-
 423
 Objectivist revisionism, 103, 298,
 304-324

Objectivity, 429-430
 Observation reports, 37, 41
 accepting as evidence, 399-400
 making, 399-400
 using as inputs in a program
 for routine decision making, 406-
 408
 for routine expansion, 399-400,
 413-419
 Optimal with respect to expected
 utility, 95-96
 Optimism-pessimism criteria, 145
 admissible (*OP*-admissible), 148
 and spreads in the odds, 154-156
 optimal (*OP*-optimal), 148
 pair, 145. *See also* Focal pair
 Ordering options with respect to
 strength, 138-143

P-admissible option, 97-98, 137-139,
 144
 Paradoxes of confirmation, 26-27
 Parsons, C., xvi
 Pearson, E. S., 37, 87, 90, 102, 156-
 157, 294-296, 326, 404-409, 415,
 419, 422-423
 Peirce, C. S., 14-18, 19, 21, 22-23,
 42, 70-72, 258, 422
 Perrin, J., 70
 Personalism,
 intemperate, 87, 91, 93, 298, 426-
 427
 tempered, 93, 298
P-function, 82
 Pivotal function, 329, 336
 invertible, 331
 smoothly, 336
 irrelevance-allowing, 332, 339
 Plato, 171
 Popper, K. R., 18-25, 31, 32, 58,
 67, 70-72, 219n, 427

Possibility
 economic, 241
 logical, 3, 242-243
 metaphysical, 190-191
 objective. *See* Ability
 physical, 241, 242
 relevant, 4, 112-118, 222-225
 serious, 2-5, 222-225, 236-237
 and ability, 244-246
 double standard of, 16-18, 72
 technological, 241, 242
 Potential answers, 45-46

Potential surprise, 183
P-predicate, 239
 Probability
 conditional, 104, 112-118, 217-219
 as confirmational commitment,
 218-219
 countably additive, 77, 104-105,
 118, 121, 122-123, 125-131, 280-
 287
 and fiducial inference, 336, 340,
 360-361
 credal, 3-4, 25, 48, 52, 234-236
 de dicto as non-truth-value-
 bearing, 185-189
 epistemological, 375
 expectation-determining, 48, 49,
 52, 104
 finitely additive, 76
 improper, 126-127, 284-287, 360
 information-determining, 48, 49-50
 metaphysical, 188-189
 pignic, 153, 213
 posterior, 84, 228, 289
 standardized, 308-309
 prior, 85, 228, 289
 standardized, 308-309
 as a science, 260
 statistical. *See* Chance
 unbiased prior, 306
 strongly unbiased prior, 306-307
 unconditional, 76
 Proliferation of rivals to settled
 assumptions, 69
 Psychologism, 67, 424-430
 Psychology, 11-12

Q-function, 76-79
 logically permissible, 85, 216
 seriously permissible, 89
Q-independence, 107
 Quine, W. V., 22, 44, 71

Raiffa, H., 145-146, 209
 Ramsey, F. P., 109, 112, 114, 116,
 210
 Random designators, 396
 Randomization, 300-302, 396
 Random membership, 392-398
 Random selection, 253, 392-398
 Ranking options with respect to
 conditional expected utility
 (RCEU), 114
 qualified (QRCEU), 116-118, 218

modified qualified (MQRCEU),
 117-118, 218
 Ranking options with respect to
 expected utility (REU), 95-96,
 106, 108-112, 114
 Regret, 147
 admissible, 147
 and spread in the odds, 154-156
 optimal, 147
 Reichenbach, H., 87, 373-374, 382
 Renyi, A., 129-130
 Replacement, 25-26, 33, 58, 63-65
 and avoidance of error, 63-64, 70-
 71
 and question begging, 65
 Residual shift, 25, 33
 Revisionism, 93
 Revolution, 63, 67-69
 Risk function, 409n
 Rule for ties, 132-135
 extended, 135-137

Sample space, 236, 248-249
 Savage, L. J., 78, 86, 89, 145n,
 149n, 200n, 210
 Schick, F., xv, 79, 199
 Security, 148, 408-413
 admissibility (*S*-admissibility), 98,
 148, 211
 and spreads in the odds, 155-156
 optimality (*S*-optimality), 148
 Seidenfeld, T., xvi, 102, 117n, 340n,
 341-342, 363-367
 Self-knowledge, 12
 Sen, A. K., 145-146n
 Shackleton, G. L. S., 146n
 Shimony, A., 56, 92-93, 99, 110,
 112, 114, 116, 117, 318, 428
 σ -finite measures, 128-131, 286, 360
 Simplicity, 43-45
 Smith, C. A. B., xvi, 79, 152-156,
 199, 200n, 210-214, 375, 387n
 Sociology, 11-12
 Spielman, S., xvi, 207, 208, 304
 Statistical model, 235-236
 Stein, H., xvi, 117n
 Stochastic irrelevance, 253-254
 and determination of likelihoods,
 298-304
 and inverse inference, 303-304
 Stokes, G. G., 63
 Stone, M., 130, 284, 287, 361
 Sufficiency, 229, 230-233

- Suppes, P., 204, 205-207, 208
 Support and its function in inquiry, 344
 Suspension of judgment, 52-53, 62, 93, 97, 133-134, 185-186, 214-215, 429
 Svedberg, T., 70
 Synthetic reasoning. *See* Corrigibility
- Tarski, A., 9, 22
 Total knowledge requirement, 80, 216
 Trade-off between risk of error and informational value, 39, 40-42, 50-51, 56
 Trials, 236
 Truth, 8-9
 avoidance of error as a desideratum in the proximate aims of inquiry, 20-23, 34, 43-45, 50-51
 equated with knowledge, 21
 myopia with respect to, 70-72
 relativized to persons and times, 21-22
 the true complete story as the ultimate goal of inquiry, 19-20, 22-24, 70-72
 avoidance of error and contraction, 59
 Tukey, J., 337
 Tversky, A., 204
- Ullian, J., 44
 Ultimate partition, 45-46, 49, 132, 181-182
u-functions, 94-95
 permissible vs. possible, 177-179
 Urcorpus (*UK*), 7, 12-13
- Valuation, rational, 75-76, 94-95
 Valuational
 closure under linear transformation, 94, 168
 consistency, 94, 168
 convexity, 94, 168, 173-176
 uniqueness, 94, 168
 Value conflict, 94-95, 168-173
 in epistemic value, 180-182
 vs. ignorance, 177-179
 Venn, G., 258
 Von Mises, R., 258-261
 Von Neumann, 163
 Wald, A., 102, 145n, 156-157, 160, 162, 163, 408-410, 422-423
 Weight of argument, 183
WU-option, 139