

OXFORD



The Lipsey Lectures

Analogies and Theories

Formal Models of Reasoning

Itzhak Gilboa,
Larry Samuelson,
and David Schmeidler

Analogies and Theories

The Lipsey Lectures

The Lipsey Lectures offer a forum for leading scholars to reflect upon their research. Lipsey lecturers, chosen from among professional economists approaching the height of their careers, will have recently made key contributions at the frontier of any field of theoretical or applied economics. The emphasis is on novelty, originality, and relevance to an understanding of the modern world. It is expected, therefore, that each volume in the series will become a core source for graduate students and an inspiration for further research.

The lecture series is named after Richard G. Lipsey, the founding professor of economics at the University of Essex. At Essex, Professor Lipsey instilled a commitment to explore challenging issues in applied economics, grounded in formal economic theory, the predictions of which were to be subjected to rigorous testing, thereby illuminating important policy debates. This approach remains central to economic research at Essex and an inspiration for members of the Department of Economics. In recognition of Richard Lipsey's early vision for the Department, and in continued pursuit of its mission of academic excellence, the Department of Economics is pleased to organize the lecture series, with support from Oxford University Press.

Analogies and Theories

Formal Models of Reasoning

Itzhak Gilboa, Larry Samuelson,
and David Schmeidler

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Itzhak Gilboa, Larry Samuelson, and David Schmeidler 2015

The moral rights of the authors have been asserted

First Edition published in 2015

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data
Data available

Library of Congress Control Number: 2014956892

ISBN 978-0-19-873802-2

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

Acknowledgments

We are grateful to many people for comments and references. Among them are Daron Acemoglu, Joe Altonji, Dirk Bergemann, Ken Binmore, Yoav Binyamini, Didier Dubois, Eddie Dekel, Drew Fudenberg, John Geanakoplos, Brian Hill, Bruno Jullien, Edi Karni, Simon Kasif, Daniel Lehmann, Sujoy Mukerji, Roger Myerson, Klaus Nehring, George Mailath, Arik Roginsky, Ariel Rubinstein, Lidror Troyanski, Peter Wakker, and Peyton Young. Special thanks are due to Alfredo di Tillio, Gabrielle Gayer, Eva Gilboa-Schechtman, Offer Lieberman, Andrew Postlewaite, and Dov Samet for many discussions that partly motivated and greatly influenced this project. Finally, we are indebted to Rossella Argenziano and Jayant Ganguli for suggesting the book project for us and for many comments along the way.

We thank the publishers of the papers included herein, The Econometric Society, Elsevier, and Springer, for the right to reprint the papers in this collection. (Gilboa and Schmeidler, "Inductive Inference: An Axiomatic Approach" *Econometrica*, 71 (2003); Gilboa and Samuelson, "Subjectivity in Inductive Inference", *Theoretical Economics*, 7, (2012); Gilboa, Samuelson, and Schmeidler, "Dynamics of Inductive Inference in a Unified Model", *Journal of Economic Theory*, 148 (2013); Gayer and Gilboa, "Analogies and Theories: The Role of Simplicity and the Emergence of Norms", *Games and Economic Behavior*, 83 (2014); Di Tillio, Gilboa and Samuelson, "The Predictive Role of Counterfactuals", *Theory and Decision*, 74 (2013) reprinted with kind permission from Springer Science+Business Media B.V.) We also gratefully acknowledge financial support from the European Research Council (Gilboa, Grant no. 269754), Israel Science Foundation (Gilboa and Schmeidler, Grants nos. 975/03, 396/10, and 204/13), the National Science Foundation (Samuelson, Grants nos. SES-0549946 and SES-0850263), The AXA Chair for Decision Sciences (Gilboa), the Chair for Economic and Decision Theory and the Foerder Institute for Research in Economics (Gilboa).

Contents

1. Introduction	1
1.1 Scope	1
1.2 Motivation	5
1.3 Overview	7
1.4 Future Directions	11
1.5 References	13
2. Inductive Inference: An Axiomatic Approach	17
2.1 Introduction	17
2.2 Model and Result	21
2.3 Related Statistical Methods	24
2.4 Discussion of the Axioms	27
2.5 Other Interpretations	31
2.6 Appendix: Proofs	32
2.7 References	46
3. Subjectivity in Inductive Inference	49
3.1 Introduction	49
3.2 The Model	52
3.3 Deterministic Data Processes: Subjectivity in Inductive Inference	56
3.4 Random Data Generating Processes: Likelihood Tradeoffs	63
3.5 Discussion	74
3.6 Appendix: Proofs	78
3.7 References	84
4. Dynamics of Inductive Inference in a Unified Framework	87
4.1 Introduction	87
4.2 The Framework	90
4.3 Special Cases	96
4.4 Dynamics of Reasoning Methods	104
4.5 Concluding Remarks	117
4.6 Appendix A: Proofs	119
4.7 Appendix B: Belief Functions	121
4.8 References	128

Contents

5. Analogies and Theories: The Role of Simplicity and the Emergence of Norms	131
5.1 Introduction	131
5.2 Framework	136
5.3 Exogenous Process	143
5.4 Endogenous Process	148
5.5 Variants	150
5.6 Appendix: Proofs	155
5.7 References	161
6. The Predictive Role of Counterfactuals	163
6.1 Introduction	163
6.2 The Framework	168
6.3 Counterfactual Predictions	173
6.4 Discussion	176
6.5 References	179
<i>Index</i>	181

1

Introduction

1.1 Scope

This book deals with some formal models of reasoning used for inductive inference, broadly understood to encompass various ways in which past observations can be used to generate predictions about future eventualities. The main focus of the book are two modes of reasoning and the interaction between them. The first, more basic, is *case-based*,¹ and it refers to prediction by analogies, that is, by the eventualities observed in similar past cases. The second is *rule-based*, referring to processes where observations are used to learn which general rules, or theories, are more likely to hold, and should be used for prediction. A special emphasis is put on a model that unifies these modes of reasoning and allows the analysis of the dynamics between them.

Parts of the book might hopefully be of interest to statisticians, psychologists, philosophers, and cognitive scientists. Its main readership, however, consists of researchers in economic theory who model the behavior of economic agents. Some readers might wonder why economic theorists should be interested in modes of reasoning; others might wonder why the answer to this question isn't obvious. We devote the next section to these motivational issues. It might be useful first to delineate the scope of the present project more clearly by comparing it with the emphasis put on similar questions in fellow disciplines.

1.1.1 Statistics

The use of past observations for predicting future ones is the bread and butter of statistics. Is this, then, a book about statistics, and what can it add to existing knowledge in statistics?

¹ The term "case-based reasoning" is due to Schank (1986) and Schank and Riesbeck (1989). As used here, however, it refers to reasoning by similarity, dating back to Hume (1748) at the latest.

While our analysis touches upon statistical questions and methods at various points, most of the questions we deal with do not belong to statistics as the term is usually understood. Our main interest is in situations where statistics typically fails to provide well-established methods for generating predictions, whether deterministic or probabilistic. We implicitly assume that, when statistical analysis offers reliable, agreed-upon predictions, rational economic agents will use them. However, many problems that economic agents face involve uncertainties over which statistics is silent. For example, statistical models typically do not attempt to predict wars or revolutions; their success in predicting financial crises is also limited. Yet such events cannot be ignored, as they have direct and non-negligible impact on economic agents' lives and decisions. At the personal level, agents might also find that some of the weightiest decisions in their lives, involving the choice of career paths, partners, or children, raise uncertainties that are beyond the realm of statistics.

In light of the above, it is interesting that the two modes of reasoning we discuss, which originated in philosophy and psychology, do have close parallels within statistics. Case-based reasoning bears a great deal of similarity to non-parametric methods such as kernel classification, kernel probabilities, and nearest-neighbor methods (see Royall, 1966, Fix and Hodges, 1951–2, Cover and Hart, 1967). Rule-based reasoning is closer in spirit to parametric methods, selecting theories based on criteria such as maximum likelihood as well as information criteria (such as the Akaike Information Criterion, Akaike, 1974) and using them for generating predictions. Case-based reasoning and kernel methods are more likely to be used when one doesn't have a clear idea about the underlying structure of the data generating process; rule-based reasoning and likelihood-based methods are better equipped to deal with situations where the general structure of the process is known. Viewed thus, one may consider this book as dealing with (i) generalizations of non-parametric and parametric statistical models to deal with abstract problems where numerical data do not lend themselves to rigorous statistical analysis; and (ii) ways to combine these modes of reasoning.

It is important to emphasize that our interest is in modeling the way people think, or should think. Methods that were developed in statistics or machine learning that may prove very successful in certain problems are of interest to us only to the extent that they can also be viewed as models of human reasoning, and especially of reasoning in the type of less structured problems mentioned above.

1.1.2 *Psychology*

If this book attempts to model human reasoning, isn't it squarely within the realm of psychology? The answer is negative for several reasons. First,

following the path-breaking contributions of Daniel Kahneman and Amos Tversky, psychological research puts substantial emphasis on “heuristics and biases”, that is, on judgment and decision making that are erroneous and that clearly deviate from standards of rationality. There is great value in identifying these biases, correcting them when possible and accepting them when not. However, our focus is not on situations where people are clearly mistaken, in the sense that they can be convinced that they have been reasoning in a faulty way. Instead, we deal with two modes of reasoning that are not irrational by any reasonable definition of rationality: thinking by analogies and by general theories. Not only are these modes of reasoning old and respectable, they have appeared in statistical analysis, as mentioned above. Thus, while our project is mostly descriptive in nature, trying to describe how people think, it is not far from a normative interpretation, as it focuses on modes of reasoning that are not clearly mistaken.

Another difference between our analysis and psychological research is that we view our project not as a goal in itself, but as part of the foundations of economics. Our main interest is not to capture a given phenomenon about human reasoning, but to suggest ways in which economic theorists might usefully model the reasoning of economic agents. With this goal in mind, we seek generality at the expense of accuracy more than would a psychologist. We are also primarily interested in mathematical results that convey general messages. In contrast to the dominant approach in psychology, we are not interested in accurately describing specific phenomena within a well-defined field of knowledge. Rather, we are interested in convincing fellow economists which paradigms should be used for understanding the phenomena of interest.

1.1.3 *Philosophy*

How people think, and even more so, how people should think, are questions that often lead to philosophical analysis. More specifically, how people should be learning from the past about the future has been viewed as a clearly philosophical problem, to which important contributions were made by thinkers who are considered to be primarily philosophers (such as David Hume, Charles Peirce, and Nelson Goodman, to mention but a few). As in other questions, whereas psychology tends to take a descriptive approach, focusing on actual human reasoning and often on its faults and flaws, philosophy puts a greater emphasis on normative questions. Given that our main interest also has a more normative flavor than does mainstream psychological research, it stands to reason that our questions would have close parallels within philosophy.

There are some key differences in focus between our analysis and the philosophical approach. First, philosophers seem to be seeking a higher degree of accuracy than we require. As economic theorists, we are trained to seek and are used to finding value in definitions and in formal models that are not always very accurate, and that have a vague claim to be generalizable without a specific delineation of their scope of applicability. (See Gilboa, Postlewaite, Samuelson, and Schmeidler, 2013, where we attempt to model one way in which economists sometimes view their theoretical models.) Thus, while philosophers might be shaken by a paradox, as would a scientist be shaken by an empirical refutation of a theory, we would be more willing to accept the paradox or the counter-example as an interesting case that should be registered, but not necessarily as a fatal blow to the usefulness of the model. The willingness to accept models that are imperfect should presumably pay off in the results that such models may offer. Our analysis thus puts its main emphasis on mathematical results that seem to be suggesting general insights.

Another distinction between our analysis and mainstream analytical philosophy is that the latter seems to be focusing on rule-based reasoning almost entirely. In fact we are not aware of any formal, mathematical models of case-based reasoning within philosophy, perhaps because this mode of reasoning is not considered to be fully rational. We maintain that there are problems of interest in which one has too little information to develop theories and select among them in an objective way. In such problems, it might be the case that the most rational thing to do is to reason by analogies. Hence we start off with the assumption that both rule-based and case-based reasoning have a legitimate claim to be “rational” modes of reasoning, and seek models that capture both, ideally simultaneously.

1.1.4 Conclusion

There are other fields in which inductive inference is studied. Artificial intelligence, relying on philosophy, psychology, and computer science, offers models of human reasoning in general and of induction in particular. Machine learning, a field closer to statistics, also deals with the same basic fundamental question of inductive inference. Thus, it is not surprising that the ideas discussed in the sequel have close counterparts in statistics, machine learning, psychology, artificial intelligence, philosophy linguistics, and so on.

The main contribution of this work is the formal modeling of arguments in a way that allows their mathematical analysis, with an emphasis on the ability to compare case-based and rule-based reasoning. The mathematical analysis serves a mostly rhetorical purpose: pointing out to economists

strengths and weaknesses of formal models of reasoning that they may be using in their own modeling of economic phenomena. With this goal in mind, we seek insights that appear to be generally robust, even if not necessarily perfectly accurate. We hope that the mathematical analysis reveals some properties of models that are not entirely obvious a priori, and may thereby be of help to economists in their modeling.

1.2 Motivation

Economics studies economic phenomena such as production and consumption, growth and unemployment, buying and selling, and so forth. All of these phenomena relate to human activities, or decision making. It might therefore seem very natural that we would be interested in human reasoning: presumably if we knew how people reason, we would know how they make decisions, and, as a result, which economic phenomena to expect.

This view is also consistent with a reductionist approach, suggesting that economics should be based on psychology: just as it is argued biology can be (in principle) reduced to chemistry, economics can be (in principle) reduced to psychology. From this point of view, it would seem very natural that economists would be interested in the way people think and perform inductive inference.

Economists have not found this conclusion obvious. First, the alleged reduction of one scientific discipline to another seldom implies that all questions of the latter should be of interest to the former. Chemistry need not be interested in high-energy physics, and biologists may be ignorant of the chemistry of polymers. Second, psychology has not reached the same level of success of quantitative predictions as have the “exact” sciences, and thus it may seem less promising as a basis for economics as would, say, physics be for chemistry. And, perhaps more importantly, in the beginning of the twentieth century the scientific nature of psychology was questioned. While the philosophy of science was dominated by the Received View of logical positivism (Carnap, 1923), and later by Popper’s (1934) thought, psychology was greatly influenced by Freudian psychoanalysis, famously one of the targets of Popper’s critique. Thus, psychology was not only considered to be an “inexact” or a “soft” science; many started viewing it as a non-scientific enterprise.²

In response to this background, many economists sought refuge in the logical positivist dictum that understanding how people think is unnecessary for understanding how they behave. The revealed preference paradigm came

² See Loewenstein (1988).

to the fore, suggesting that all that matters is observed behavior (see Frisch, 1926, Samuelson, 1938). Concepts such as tastes and beliefs were modeled as mathematical constructs—a utility function and a probability measure—which are defined solely by observed choices. Economists came to think that how people think, and how they form their beliefs, was, by and large, of no economic import. Or, to be precise, the beliefs of rational agents came to be modeled by probability measures which were assumed to be updated according to Bayes's rule with the arrival of new information. It became accepted that, beyond the application of Bayes's rule for obtaining conditional probabilities, no reasoning process was necessary for understanding people's choices and resulting economic phenomena.

This view of economic agents as “black boxes” that behave *as if* they were following certain procedures paralleled the rise of behaviorism in psychology (Skinner, 1938). Whereas, however, in psychology, strict behaviorism was largely discarded in favor of cognitive psychology (starting in the 1960s), in economics the “black box” approach survives to this day. (See, for instance, Gul and Pesendorfer, 2008.) Indeed, given that the subject matter of economics is people's economic activities, it is much easier to dismiss mental phenomena and cognitive processes as irrelevant to economics than it is to do so when discussing psychology. And, importantly, axiomatic treatments of people's behavior, and most notably Savage's (1954) result, convinced economists that maximizing expected utility relative to a subjective probability measure is the model of choice for descriptive and normative purposes alike. This model allows many degrees of freedom in selecting the appropriate prior belief, but beyond that leaves very little room for modeling thinking. Presumably, if we know how people behave and make economic decisions, we need not concern ourselves with the way people think.

We find this view untenable for several reasons. First, Savage's model is hardly an accurate description of people's behavior. In direct experimental tests of the axioms, a non-negligible proportion of participants end up violating some of them (see Ellsberg, 1961, and the vast literature that followed). Moreover, many people have been found to consistently violate even more basic assumptions (see Tversky and Kahneman, 1974, 1981). Further, when tested indirectly, one finds that many empirical phenomena are easier to explain using other models than they are using the subjective expected utility hypothesis. Hence, one cannot argue that economics has developed a theory of behavior that is always satisfactorily accurate for its purposes. It stands to reason that a better understanding of people's thought processes might help us figure out when Savage's theory is a reasonable model of agents' behavior, and how it can be improved when it isn't.

Second, Savage's result is a powerful rhetorical device that can be used to convince a decision maker that she would like to conform to the subjective

expected utility maximization model, or even to convince an economist that economic agents might indeed behave in accordance with this model, at least in certain domains of application. But the theorem does not provide any guidance in selecting the utility function or the prior probability involved in the model. Since tastes are inherently subjective, theoretical considerations may be of limited help in finding an appropriate utility function, whether for normative or for descriptive purposes. However, probabilities represent beliefs, and one might expect theory to provide some guidance in finding which beliefs one should entertain, or which beliefs economic agents are likely to entertain. Thus, delving into reasoning processes might be helpful in finding out which probability measures might, or should capture agents' beliefs.

Third, Savage's model follows the general logical positivistic paradigm of relating the theoretical terms of utility and probability to observable choice. But these choices often aren't observable in practice, and sometimes not even in principle. For example, in order to capture possible causal theories, one needs to construct the state space in such a way that it is theoretically impossible to observe the preference relation in its entirety. In fact, observable choices would be but a fraction of those needed to execute an axiomatic derivation. (See Gilboa and Schmeidler, 1995, and Gilboa, Postlewaite, and Schmeidler, 2009, 2012.) Hence, for many problems of interest one cannot rely on observable choice to identify agents' beliefs. On this background, studying agents' reasoning offers a viable alternative to modeling beliefs.

In sum, we believe that understanding how people think might be useful in predicting their behavior. While in principle one could imagine a theory of behavior that would be so accurate as to render redundant any theory of reasoning, we do not believe that the current theories of behavior have achieved such accuracy.

1.3 Overview

The present volume consists of six chapters, five of which have been previously published as separate papers. The first two of these deal with a single mode of reasoning each, whereas the rest employ a model that unifies them. Chapter 2³ focuses on case-based reasoning. It offers an axiomatic approach to the following problem: given a database of observations, how should different eventualities be ranked? The axiomatic derivation assumes that observations in a database may be replicated at will to generate a new database, and that it would be meaningful to pose the same problem for the

³ Gilboa and Schmeidler, 2003.

new database. For example, if the reasoner observes the outcomes of a roll of a die, and has to predict which outcome is more likely to occur on the next roll, we assume that any database consisting of finitely many past observations can be imagined, and that the reasoner should be able to respond to the ranking question given each such database. The key axiom, *combination*, roughly suggests that, should eventuality a be more likely than another eventuality b , given two disjoint databases, then a should be more likely than b also given their union. Ranking outcomes by their relative frequencies clearly satisfies this axiom: if one outcome has appeared more often than another in each of two databases, and will thus be considered more likely given each, so it will be given their union. Coupled with a few other, less fundamental assumptions, the combination axiom implies that the reasoner would be ranking alternative eventualities by an additive formula. The formula can be shown to generalize simultaneously several known techniques from statistics, such as ranking by relative frequencies, kernel estimation of density functions (Akaike, 1945), and kernel classification. Importantly, the model can also be applied to the ranking of theories given databases, where it yields an axiomatic foundation for ranking by the maximum likelihood principle.⁴ The chapter also discusses various limitations of the combination axiom. Chief among them are situations in which the reasoner engages in second-order induction, learning the similarity function to be used when performing case-to-case induction,⁵ and in learning that involves both case-to-rule induction and (rule-to-case) deduction. These limitations make it clear that, while the combination axiom is common to several different techniques of inductive inference, it by no means encompasses all forms of learning.

Chapter 3⁶ deals with rule-based reasoning. It offers a model in which a reasoner starts out with a set of theories and, after any finite history of observations, needs to select a theory. It is assumed that the reasoner has a subjective a priori ranking of the theories, for example, a “simpler than” relation. Importantly, we assume that there are countably many theories, and for each one of them there are only finitely many other theories that are ranked higher. Given a history, the reasoner rules out those theories that have been refuted by the observations, and selects a maximizer of the subjective ranking among those that have not been refuted, that is, chooses one of the simplest theories that fit the data. A key insight is that, in the absence of a subjective ranking, the reasoner would not be able to learn effectively: she would be unable to consistently choose among all possible theories that are consistent with observed history. Hence, even if the observations happen to

⁴ A sequel paper, Gilboa and Schmeidler (2010), generalizes the model to allow for an additive cost attached to a theory's log-likelihood, as in Akaike Information Criterion.

⁵ See Gilboa, Lieberman, and Schmeidler, 2006.

⁶ Gilboa and Samuelson, 2012.

fit a simple theory, the reasoner will not conclude that this theory is to be used for prediction, as there are many other competing theories that match the data just as well. By contrast, when a subjective ranking—such as simplicity—is used as an additional criterion for theory selection, the reasoner will learn simple processes: at some point all theories that are simpler than the true one (but not equivalent to it) will be refuted, and from that point on the reasoner will use the correct theory for prediction. Thus, the preference for simplicity provides an advantage in prediction of simple processes, while incurring no cost when attempting to predict complex or random processes. This preference for simplicity does not derive from cognitive limitations or the cost of computation; simplicity is simply one possible criterion that allows the reasoners to settle on the correct theory, should there be one that is simple. In a sense, the model suggests that had cognitive limitations not existed, we should have invented them.

Chapter 4⁷ offers a formal model that captures both case-based and rule-based reasoning. It is also general enough to describe Bayesian reasoning, which may be viewed as an extreme example of rule-based reasoning. The reasoner in this model is assumed to observe the unfolding of history, and, at each stage t , after observing some data, x_t , to make a single-period prediction by ranking possible outcomes in that period, y_t . The reasoner uses *conjectures*, which are simply subsets of states of the world (where each state specifies (x_t, y_t) for all t). Each conjecture is assigned a non-negative weight a priori, and after each history those conjectures that have not yet been refuted are used for prediction. As opposed to Chapter 3, here we do not assume that the reasoner selects a single “most reasonable conjecture” in each period for generating predictions; rather, all unrefuted conjectures are consulted, and their predictions are additively aggregated using their a priori weights. (The model also distinguishes between relevant and irrelevant conjectures, though the ranking of eventualities in each period is unaffected by this distinction).

The extreme case in which all weight is put on conjectures that are singletons (each consisting of a single state of the world) reduces to Bayesian reasoning: the a priori weights are then the probabilities of the states, and the exclusion of refuted conjectures boils down to Bayesian updating. The model allows, however, a large variety of rules that capture non-Bayesian reasoning: the reasoner might believe in a general theory that does not make specific predictions in each and every period, or that does not assign probabilities to the values of x_t . More surprisingly, the model allows us to capture case-based reasoning, as in kernel classification, by aggregating over appropriately defined “case-based conjectures”. Beyond providing a unified framework for these modes of reasoning, this model also allows one to ask

⁷ Gilboa, Samuelson, and Schmeidler, 2013.

how the relative weights of different forms of reasoning might change over time. We show that, if the reasoner does not know the structure of the underlying data generating process, and has to remain open-minded about all possible eventualities, she will gradually use Bayesian reasoning less, and shift to conjectures that are not as specific. The basic intuition is that, because Bayesian reasoning requires that weight of credence be specified to the level of single states of the world, this weight has to be divided among pairwise disjoint subsets of possible histories, and the number of these subsets grows exponentially fast as a function of time, t . If the reasoner does not have sharp *a priori* knowledge about the process, and hence divides the weight of credence among the subsets in a more or less unbiased way, the weight of each such subset of histories will be bounded by an exponentially decreasing function of t . By contrast, conjectures that allow for many states may be fewer, and if there are only polynomially many of them (as a function of t), their weight may become relatively higher as compared to the weight of the Bayesian conjectures. This result suggests that, due to the fact that Bayesian approach insists on quantifying *any* source of uncertainty, it might prove non-robust as compared to modes of reasoning that remain silent on many issues and risk predictions only on some.

Chapter 5⁸ uses the same framework to focus on case-based vs. rule-based reasoning. Here, the latter is understood to mean theories that make predictions (regarding y_t) at each and every period (after having observed x_t), so, in this model theories cannot “pick their fights”, as it were. They differ from Bayesian conjectures in that the latter are committed to predict not only the outcome y_t but also the data x_t . Yet, making predictions about y_t at each and every period is sufficiently demanding to partition the set of unrefuted theories after every history, and thereby to generate an exponential growth of the number of subsets of theories that may be unrefuted at time t . In this chapter it is shown that, under certain reasonable assumptions, should reality be simple, that is, described by state of the world that conforms to a single theory, the reasoner will learn it. The basic logic of this simple result is similar to that of Chapter 3: it suffices that the reasoner be open-minded to conceive of all theories and assign some weight to them. Should one of these simple theories be true, sooner or later all other theories will be refuted, and the *a priori* weight assigned to the correct theory will become relatively large. Moreover, in this chapter we also consider case-based conjectures, and show that their weight diminishes to zero. As a result, not only is the correct theory getting a high weight relative to other theories, the entire class of rule-based conjectures becomes dominant as compared to the case-based ones. That is, the reasoner would converge to be rule-based.

⁸ Gayer and Gilboa, 2014.

However, in states of the world that are not simple, that is, that cannot be described by a single theory, under some additional assumptions the converse is true: similarly to the analysis of Chapter 4, case-based reasoning would drive out rule-based reasoning. Chapter 5 also deals with situations in which the phenomenon observed is determined by people's reasoning, that is, that the process is endogenous rather than exogenous. It is shown that under endogenous processes rule-based reasoning is more likely to emerge than under exogenous ones. For example, it is more likely to observe people using general theories when predicting social norms than when predicting the weather.

Finally, Chapter 6⁹ applies the model of Chapter 4 to the analysis of counterfactual thinking. It starts with the observation that, while counterfactuals are by definition devoid of empirical content, some of them seem to be more meaningful than others. It is suggested that counterfactual reasoning is based on the conjectures that have not been refuted by actual history, h_t , applied to another history, h'_t , which is incompatible with h_t (hence counterfactual). Thus, actual history might be used to learn about general rules, and these can be applied to make predictions also in histories that are known not to be the case. This type of reasoning can make interesting predictions only when the reasoner has non-Bayesian conjectures: because each Bayesian conjecture consists of a single state of the world, a Bayesian conjecture that is unrefuted by the actual history h_t would be silent at the counterfactual history h'_t . However, general rules and analogies that are unrefuted by h_t might still have non-trivial predictions at the incompatible history h'_t . The model is also used to ask what counterfactual thinking might be useful for, and to rule out one possible answer: a rather trivial observation shows that, for an unboundedly rational reasoner, counterfactual prediction cannot enhance learning.

1.4 Future Directions

The analysis presented in this volume is very preliminary and may be extended in a variety of ways. First, in an attempt to highlight conceptual issues, we focus on simple models. For example, we assume that theories are deterministic; and that case-based reasoning takes into account only the similarity between two cases at a time. In more elaborate models, one might consider probabilistic theories, analogies that involve more than two cases, more interesting hybrids between case-based and rule-based theories, and so forth.

⁹ Di Tillio, Gilboa, and Samuelson, 2013.