Available online at www.sciencedirect.com

SciVerse ScienceDirect

Journal of Economic Theory 148 (2013) 1399–1432

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Dynamics of inductive inference in a unified framework ☆

Itzhak Gilboa [a,b,c,*,1], Larry Samuelson [d,2], David Schmeidler [a,e,f,3]

[a] *Tel-Aviv University, Israel*
[b] *HEC, Paris, France*
[c] *Cowles Foundation, Yale University, United States*
[d] *Department of Economics, Yale University, United States*
[e] *The Ohio State University, United States*
[f] *The InterDisciplinary Center in Herzliya, Israel*

Received 25 October 2011; final version received 22 May 2012; accepted 26 November 2012

Available online 14 January 2013

Abstract

We present a model of inductive inference that includes, as special cases, Bayesian reasoning, case-based reasoning, and rule-based reasoning. This unified framework allows us to examine how the various modes of inductive inference can be combined and how their relative weights change endogenously. For example, we establish conditions under which an agent who does not know the structure of the data generating process will decrease, over the course of her reasoning, the weight of credence put on Bayesian vs. non-Bayesian reasoning. We illustrate circumstances under which probabilistic models are used until an unexpected outcome occurs, whereupon the agent resorts to more basic reasoning techniques, such as case-based and rule-based reasoning, until enough data are gathered to formulate a new probabilistic model.

JEL classification: D8

Keywords: Induction; Learning; Analogies; Theories; Case-based reasoning; Rule-based reasoning

☆ We thank Daron Acemoglu, Dirk Bergemann, Eddie Dekel, Drew Fudernberg, Gabi Gayer, Offer Lieberman, George Mailath, an associate editor, and two referees for comments and suggestions.

* Corresponding author at: Tel-Aviv University, Israel.
E-mail address: tzachigilboa@gmail.com (I. Gilboa).

1 ISF Grant 396/10 and ERC Grant 269754 are gratefully acknowledged.
2 National Science Foundation grant SES-0850263 is gratefully acknowledged.
3 ISF Grant 396/10 is gratefully acknowledged.

0022-0531/$ – see front matter
http://dx.doi.org/10.1016/j.jet.2012.11.004

1. Introduction

Economic theory typically assumes that agents reason about uncertainty in a Bayesian way: they formulate prior probabilities over a state space and update them in response to new information according to Bayes' rule. This model is powerful, but does not always reflect the way that people think about uncertainty. In particular, when completely unexpected outcomes occur, people question their probabilistic models, relying on alternative reasoning techniques until perhaps developing a new probabilistic model.

For example, the New York Stock Exchange was closed for five days following the September 11, 2001 terrorist attacks on the United States. On the following Sunday, September 16, a leading economist was asked to predict the behavior of the Dow Jones Industrial Average on Monday. He did not respond by reasoning that "I used to attach (the quite small) probability ε to such attacks, and now I need only update this probability, and then apply my usual model of the stock market." Instead, there was a sense that the probabilistic model he would have used under normal circumstances was inappropriate for the present situation, and that he had to start from basics in reasoning about the future. He responded by invoking analogies to past cases in which the United States had been surprised by attack, most notably Pearl Harbor. (As it turned out, his prediction was quite accurate.)

Similarly, following the collapse of Lehman Brothers in September 2008, the head of a major investment firm confronted clients anxious to sell their assets, even assets that had already lost 90% of their value. Again, the analyst did not apply Bayes' rule to a prior that had taken into account a possible failure of Lehman Brothers. Instead, he argued that something totally unexpected had happened, and that "obviously, the models do not work." The analyst convinced his clients to hold such assets, invoking the simple rule that "an asset that has lost 90% of its value cannot lose much more." (His clients were convinced, and subsequently appreciated the advice.)

In both examples, one could, post-hoc, construct a prior probability distribution that allows the experts' reasoning to follow from Bayesian updating. However, such a description would say very little about the actual reasoning process of the agents involved, and (more importantly) would not be of much help in predicting their reasoning in the future. Our interest in this paper is in modeling the agents' actual reasoning processes, in the hope of better understanding when these processes generate probabilistic beliefs, which beliefs are likely to be formed by the agents, and how the agent might form beliefs when driven away from familiar probabilistic models. To do so, we need a model that can simultaneously describe probabilistic and non-probabilistic reasoning, as well as the dynamics by which weights shift between modes of reasoning.

We take it for granted that when statistical analysis is possible, rational agents will perform such analysis correctly. In contrast, our interest is in the way economists model agents who face problems that do not naturally lend themselves to statistical analysis. Predicting financial crises, economic growth, the outcome of elections, or the eruptions of wars and revolutions, are examples where it is difficult to define iid random variables and, more generally, where the assumptions of statistical models do not seem to be good approximations.

To set the context for our model, consider an agent who each year is called upon to predict the price of oil over the subsequent year. To keep this illustrating example simple, suppose the agent need only predict whether the average price will be higher or lower than the previous year's price. We can imagine the agent working for a hedge fund that is interested in whether it should bet for or against an increasing price.

To support her decision, the agent's research staff regularly compiles a list of data potentially relevant to the price of oil, as well as data identifying past values of the relevant variables and