



Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
<http://www.econ.upenn.edu/pier>

PIER Working Paper 08-002

“Probability and Uncertainty in Economic Modeling”
Second Version

by

Itzhak Gilboa, Andrew W. Postlewaite, and David Schmeidler

<http://ssrn.com/abstract=1088122>

Probability and Uncertainty in Economic Modeling

Itzhak Gilboa, Andrew W. Postlewaite, and David Schmeidler

Itzhak Gilboa is Professor of Economics, Tel-Aviv University, Tel Aviv, Israel, and Professor of Finance and Economics, HEC, Paris, France. Andrew Postlewaite is Harry P. Kamen Professor of Economics and Professor of Finance, University of Pennsylvania, Philadelphia, Pennsylvania. David Schmeidler is Professor of Economics, Ohio State University, Columbus, Ohio. Their e-mail addresses are <tzachigilboa@gmail.com>, <apostlew@econ.upenn.edu> and <schmeid@tau.ac.il>, respectively.

Abstract

Economic modeling assumes, for the most part, that agents are Bayesian, that is, that they entertain probabilistic beliefs, objective or subjective, regarding any event in question. We argue that the formation of such beliefs calls for a deeper examination and for explicit modeling. Models of belief formation may enhance our understanding of the probabilistic beliefs when these exist, and may also help up characterize situations in which entertaining such beliefs is neither realistic nor necessarily rational.

Since the early days of probability theory, there has been a distinction between probabilities that are given, as in a game of chance, and probabilities that are not given, but reflect a subjective degree of belief. (See Hacking, 1975, and Shafer, 1978 for historical surveys.) In economics, Knight (1921) is typically credited with the distinction between situations of “risk” and of “uncertainty”. The former designates situations in which probabilities are known, or knowable in the sense that they can be estimated from past data, calculated using the laws of probability, etc. By contrast, “uncertainty” refers to situations in which probabilities are neither known, nor can they be deduced, calculated, or estimated in an objective way. The Bayesian approach minimizes the importance of this distinction by introducing the notion of “subjective probability”. According to this approach, when objective probabilities are not known, they can be replaced by subjective ones, thus reducing problems of decision under uncertainty to decisions under risk.

The standard practice in economics when modeling situations of uncertainty is to follow the Bayesian approach and to assume that people have probabilistic beliefs over any source of uncertainty, that they update these beliefs in accordance with Bayes’s rule, and that they use these probabilistic beliefs in decision making, typically as a basis for expected utility maximization. This paradigm is an elegant and coherent way to deal with uncertainty. Yet, it is not always clear how subjective beliefs should be formed. The following two examples illustrate the problem.

Example 1: Ann is an admission officer for an economics graduate program. Every year she reviews a large number of applications to assess the probability that various candidates are likely to succeed in the program if admitted. Faced with a particular file, how should she assign a probability to the candidate’s success?

Example 2: Bob is a high school graduate considering joining the military. The advantages of a military career are relatively clear to him. The possible costs, however, are subject to uncertainty. Specifically, Bob realizes that, if the United States engages in a war in the next few years, he is likely to be stationed overseas, be involved in combat, and risk his life. To make a rational decision, Bob attempts to assess the probability of such a war.

Both Ann and Bob would prefer to have probabilistic assessments that are “objective” or even “scientific”. However, there are no agreed-upon methods for assigning probabilities to the events that concern them. Laplace’s “principle of indifference” – also known as the “principle of insufficient reason” – suggests assigning equal probabilities to all possible eventualities, say, 50%-50% chance of success vs. failure, or of war vs. peace. This approach is clearly not suitable in these problems, where sufficient information is available to break the symmetry between the two possible

outcomes. Indeed, Ann and Bob should know better than to assign a default of 50% to each outcome.

However, the existing empirical evidence is of limited use. When Ann considers a particular candidate, she does not wish to rely on the overall percentage of students who have succeeded, because these students differ in a variety of variables, and some students' performance appears more relevant than others. It would be natural for Ann to run a regression using observable characteristics of past applicants on some measure of their success in the program. But quantifiable characteristics may leave much to Ann's intuition. Moreover, the choice of the variables for the regression model, as well the details of the statistical procedure, introduces a certain degree of subjectivity.

Bob's belief formation problem is even more difficult than Ann's. He might attempt to employ an empirical frequency approach for his problem. He would then consider a database consisting of conflict situations, and calculate the percentage of these situations that resulted in war. This relative frequency might be taken as a proxy for the probability of a war occurring in the next few years. But Bob will be considering a database of conflict situations that differ from each other in a variety of variables. Some of them will be recent; others will not. Some will involve similar countries, and some won't. Taking all recorded conflicts into account would be unreasonable, but taking only "identical" situations would result in an empty comparison set. Moreover, regression analysis seems even more problematic in Bob's case than in Ann's. The resolutions of different conflicts throughout history are causally interdependent in intricate ways that are not precisely known, and that cannot be captured by simple autoregression models. Thus, neither empirical frequencies nor regression analysis offer a clear procedure for the generation of beliefs.

In many economic problems of interest it is not clear how one should define probabilities. Probabilities are actually "given" in only very restricted situations such as state lotteries or casino games. In other situations, such as insurance problems, probabilities can be reasonably approximated by relative frequencies of comparable instances computed from publicly available data. But in a vast range of economic problems, probabilities are neither explicitly given nor can they be approximated by relative frequencies or regression analysis. The goal of this paper is to ask how probabilities should be defined in these situations, and, if probabilities cannot be defined in a satisfactory way, how beliefs should be modeled.

We begin with a discussion of subjective probability, which is the standard approach to problems involving uncertainty, and which relies on well-known axiomatic foundations. We will argue, however, that the axioms underlying subjective probability are in some ways too restrictive, and in other ways too general: on the one hand, we hold that rational decision-makers may violate the axioms, while, on the other hand, the

axioms do not restrict probabilistic beliefs in any way.¹ With this background in place, we consider the implications of these limitations of the Bayesian approach to economic modeling. On the one hand, because the Bayesian model is too restrictive, one may wish to consider alternative models. We discuss such models, which account for a wider set of phenomena than the Bayesian one, and offer a few examples of economic models that make use of such approaches. On the other hand, because the Bayesian model does not say enough about the generation of beliefs, we point to a research agenda that will develop formal, explicit theories of the belief formation process. Such theories may help both to refine the predictions of economic models when agents are Bayesian, and to improve these predictions when they are not.

Subjective Probabilities

In what has become the classical theory of consumer decision-making, von Neumann and Morgenstern (1944) offered a set of axioms for maximization of expected utility. Their axioms are stated in the language of probabilities, and their theory is therefore restricted to situations risk, rather than uncertainty. It does not tell us where probabilities might emerge from if they are not given, and it does not support expected utility maximization if uncertainty is not quantified. Savage's (1954) main contribution was to extend the expected utility paradigm to situations where objective probabilities may not exist. He showed that axioms on consistency of choice in the face of uncertainty imply that the decision-maker behaves as if he had a subjective probability, with respect to which he wished to maximize expected utility. Deriving probability and utility simultaneously from observed choices, Savage has provided the most compelling justification of the Bayesian paradigm: he showed that in order to make coherent choices, individuals should behave as though they had subjective probabilities, even when objective probabilities cannot be defined. Savage's approach merits a more detailed discussion.

Savage's Axiomatic Approach

Savage's (1954) axiomatization considers observable choices between pairs of uncertain acts and rules out patterns of choices that do not seem "sensible".² To understand the

¹ This statement refers to the standard Bayesian decision model. Many economic models augment this model with very strong assumptions about beliefs, such as an assumption of rational expectations or an assumption that there is a common prior in multi-agent problems. However, the axiomatic foundations of individual beliefs, which are our focus here, do not imply these additional assumptions. In this paper we deal with individual (as opposed to interactive) decision making.

² There are other axiomatizations of subjective probability, coupled with the principle of expected utility maximization: for example, Ramsey (1931), de Finetti (1937), and Anscombe-Aumann (1963). We focus

axioms, consider bets on a horse race. Suppose you were offered the following gambles: “If Horse A wins the race you get a trip to Paris (otherwise you get nothing)” or “If Horse B wins the race you get a trip to London (otherwise you get nothing)”. Choices between such gambles reflect both the desirability of the outcomes and the probability of the events in question. For instance, if you chose the first gamble, an outside observer might suspect that you thought that it was more likely that Horse A would win the race than Horse B, or that you preferred Paris to London, or that some combination of your beliefs about the likely winner of the race and your preferences over the two cities led to the observed choice. Savage suggested axioms, stated in the language of preferences between such gambles, that suffice for the identification of *both* a utility function and a probability measure that jointly characterize the decision-maker through representation of the his choices by maximization of his subjective expected utility.

Savage postulated four conceptually important axioms.³ The first is the classical assumption, familiar from consumer theory, that preferences are complete and transitive. *Completeness* states that for any two bets, the decision-maker can say which is (weakly) preferred, that is, at least as good as the other. Offered bets “If Horse A or B wins the race you get a trip to Paris (otherwise nothing)” and “If Horse C does *not* win the race you get a trip to London (otherwise nothing)”, the person has well-defined preferences between them (including the possibility of indifference). *Transitivity* requires that (weak) preference for gamble A over B and for B over C results in (weak) preference of A over C.

Two additional axioms deal with the separation of tastes from beliefs. These axioms are quite restrictive in particular applications that involve uncertainty about one’s health or survival, where preferences can alter depending on the state in which the preferences occur, but these concerns are not the main focus of this paper and we therefore do not discuss these axioms here. For a discussion and critique of these axioms, see Drèze, (1961), Karni, Schmeidler, and Vind, (1983), and Karni (1993, 1996, 1999, 2003).

A fourth axiom, which is crucial for the present discussion, is the *Sure Thing Principle*. To illustrate this axiom, consider the following gambles. Gamble G1: “If Horse A wins the race you will get a trip to Paris and if Horse A does *not* win you will get a trip to Philadelphia”; Gamble G2: “If Horse A wins the race you will get a trip to London, and if Horse A does *not* win you will get a trip to Philadelphia”. The two gambles are identical if Horse A does not win the race, but offer a choice between Paris and London if Horse A does win. Consider also two other gambles in which the prizes if Horse A wins are the same, but there is a different consolation prize if Horse A doesn’t win. Specifically, let Gamble G3 be “If Horse A wins the race you will get a trip to Paris and if Horse A does *not* win you will get a trip to Montreal”; Gamble G4: “If Horse A wins the race you will get a trip to London, and if Horse A does *not* win you will get a trip to Montreal”. G1

on Savage’s axiomatization because it is widely perceived to be the most satisfactory from a conceptual viewpoint.

³ Savage also needed three additional axioms, which can be viewed as technical. They guarantee notions of continuity and rule out trivial cases. These axioms will not be discussed here.

and G2 differ only in the case that Horse A wins, and then the question is whether the decision-maker prefers Paris to London. Similarly, G3 and G4 differ only in the case that Horse A wins, and again, the question is whether the decision-maker prefers Paris to London. The Sure Thing Principle requires that the decision-maker prefer G1 to G2 if and only if the decision-maker prefers G3 to G4.

These and the other Savage (1954) axioms seem eminently reasonable. Savage's theorem states that, if a decision-maker's choices are coherent, in the sense that they satisfy the axioms, then these choices are equivalent to the maximization of expected utility with respect to a subjective probability measure. That is, the decision-maker behaves as though he had a probability distribution over the states of the world (specifying which horse wins the race in the example above) and a utility function over the outcomes (trips to Paris, London, Philadelphia, and Montreal above), and the decision-maker maximizes the sum of the utilities of the outcomes weighted by the probabilities that the outcomes will occur.

Why Savage's Result Matters

Savage's (1954) theorem that a person behaves as though he maximizes expected utility is important for several reasons. First, it can be interpreted normatively: to the extent that the axioms appear reasonable, so does expected utility theory. Thus, if Ann or Bob were to ask us for a recommended course of action in their respective decision problems, we might start by asking them whether they would like to make decisions consistent with these axioms. Assume that Ann considers Savage's axioms and says, "Yes, this is the kind of decision-maker I'd like to be" or even, "Now that you explained the axioms to me, I would be embarrassed to be caught violating these axioms." We can then quote Savage's theorem and say, "Well, then, you must behave as if you were maximizing the expectation of a certain utility with respect to a certain probability measure. It would, perhaps, be easier for you to think directly in terms of utility and probability and, once you chose these functions, simply to follow expected utility maximization."

Second, the axiomatization is useful for descriptive purposes. It delineates the scope of observed phenomena that are consistent with subjective expected utility theory – namely, the claim that people maximize expected utility relative to a subjective probability measure. As such, the axiomatization may also help in testing subjective expected utility theory. Because direct empirical tests of the theory may be fraught with identification problems (in real world situations, isolating effects of different subjective probabilities or different utilities is likely to be quite difficult), one may wish to test Savage's axioms in simple choice situation in the laboratory or in mind experiments. To the extent that the axioms appear valid in such experiments, one might be convinced that they are also valid in real choice situations, and, therefore, that subjective expected utility maximization is a good model of the way people make decisions in reality.

To illustrate this point, suppose that economist A, when analyzing the military career choice that Bob and other young men make, assumes that they make decisions by

maximizing their subjective expected utility. Economist B is skeptical that this is the right model to use. Neither economist has direct access to the decision processes of the young men in question. Moreover, the economists do not have sufficient data on the choices made by these men to test whether they generally are expected utility maximizers. If A were to suggest the literal interpretation of expected utility theory, namely, that Bob actually calculates products of utilities and probabilities, B would find A's theory bizarre. But suppose that A goes over Savage's axioms, and asks B whether it is plausible that the decisions Bob would make, given various choice situations, would be in accordance with the axioms and that B finds these consistency requirements reasonable. Then A may quote Savage's theorem, convincing B to accept expected utility theory as a descriptive theory of the decisions that will be made to the same degree that B accepts the axioms as description of behavior. This is *not* an argument that A will be able to convince B that maximizing expected utility with respect to a subjective probability distribution is a good description of the *process* by which the young men reach their decisions, only that their decisions are the same *as though* they did so.

Third, Savage's axiomatization of subjective expected utility maximization can also facilitate the determination of the subjective probability of various events by focusing on simple trade-offs. Suppose, for example, that Ann would like to elicit her subjective probability for the event "Candidate X will graduate successfully from the program." She might ask herself questions such as "Do I prefer to bet on X graduating successfully or on another candidate, Y?" "Am I willing to bet on X graduating vs. failing at odds 1:2?" Such preference questions may have a more palpable meaning to Ann than the question, "What is the precise probability p that X will graduate?" Yet, if Ann satisfied Savage's axioms, a set of simple binary comparisons like these will identify a unique p that can be defined as her subjective probability for the event in question.

Difficulties in Savage's Approach to Rationality

Savage's axioms are often considered to be the behavioral definition of rationality. If Ann or Bob ask us what will be a rational decision for them to make, we should point out to them that, if they accept the axioms, then there must be probabilities and utilities that represent their choices via expected utility maximization. They may then ask us, how they can figure out which probability and which utility are "theirs". Here we can resort to Savage's axiomatization again, and tell them that they should examine their own preferences and elicit their implicit beliefs. This approach, however, faces several difficulties.

First, Ann and Bob might find themselves expressing preferences that are in contradiction to one or more of the axioms. For example, assume that Ann is considering files of two candidates. Candidate X comes from a college Ann knows well. She has seen many similar students, and she observes that about 60 percent of them graduated successfully. Candidate Y comes from a foreign country. Ann has no experience with students from Y's college, or, in fact, with anyone from Y's country. Out of ignorance she might assign to this candidate a success probability that is the overall success rate for all students in the

program. Assume that this general success rate is also 60 percent. Yet, Ann knows that this number, 60 percent, was assigned almost as a default. By contrast, the 60 percent probability assigned to the success of candidate X is based on a significant amount of information. The two probabilities, though equal, “feel” different. More concretely, we should not be surprised if Ann is more willing to bet on candidate X’s success than on Y’s. Typically, Ann might feel safer with a bet whose distribution is known, rather than with one whose distribution is not known.

The Bayesian approach, logically necessitated by Savage’s axioms, fails to distinguish between probabilities based on data and probabilities that result from a default rule (which is another way of saying “ignorance”). Consider bets on two coins, one which was extensively tested and was found to be fair, and another about which nothing is known. The outcome of a toss of the first coin will be assigned a 50-50 distribution due to “hard” evidence. The outcome of a toss of the second coin will be assigned the same distribution in accord with Laplace’s principle of indifference. But as Schmeidler (1989) argues, the two distributions feel different, and, as a result, our willingness to bet on them need not be the same. In a classic experiment, Ellsberg (1961) has shown that people often express preferences for bets with known probabilities over bets with unknown probabilities.⁴

Another fundamental difficulty with the descriptive interpretation of the axioms underlying expected utility theory has to do with the completeness axiom, namely, the assumption that the decision-maker has preferences between any two uncertain acts (Shafer, 1986; Bewley, 2002; Gilboa, Postlewaite, and Schmeidler, 2006). Ann and Bob may find that, for many pairs of acts, they simply do not have well-defined preferences. Specifically, if we were to ask Bob if he preferred to join the military or not, his reply would likely be: “That is precisely what I am trying to find out.” Likewise, Ann’s choice between different candidates is the decision problem for which she is interested in probabilities in the first place.

The completeness axiom is a standard assumption in consumer theory. Indeed, when the outcomes of various choices are certain and known to a consumer, this axiom is rather innocuous. If Dan is offered a choice between a bowl of chocolate ice cream and a bowl of vanilla ice cream, Dan is likely to choose a particular flavor with no hesitation. He doesn’t need to make any calculations about the options. The consumer has well-defined preferences, which are also accessible to him through introspection. Correspondingly, if he prefers chocolate to vanilla, no outsider can convince him that his choice is incorrect, and that he actually prefers vanilla to chocolate.

⁴ There are other reasons why Savage’s (1954) axioms might be violated by observed behavior. Some involve general critiques of the rational choice paradigm, such as framing effects, gain-loss asymmetry and other phenomena documented in Tversky and Kahneman (1974). Violations of transitivity were observed by Lichtenstein and Slovic (1971). Other problems are specific to the expected utility model, such as state-dependent preferences briefly mentioned above.

But in the presence of uncertainty, whether about objective outcomes or about one's subjective experience, completeness of preferences is a less compelling assumption. Assume that Carol is taking a new job, and she is offered either one pension plan with defined benefits that depend in a complex way on wages and years of service, or a plan with defined contributions whose eventual pension payments will depend on amount contributed, return on investment, and choices over types of payouts made at retirement. Carol must make a choice, as must Dan, who has to choose what flavor of ice cream to have. But Carol's choice is very different from Dan's. Carol is likely to have no *a priori* preferences between the plans, unlike the choice between chocolate and vanilla. In Carol's case, fact-finding and reasoning must precede preferences. All that Carol can say *a priori* is that, at any point in the future, and at every realization of uncertainty, say, about her health, she prefers more money to less. But these preferences, which precede reasoning, do not suffice to determine complicated choices between uncertain prospects. In particular, the two pension plans would have to be *analyzed* for (or by) Carol to determine her preferences between them. The analysis of the two plans calls for the assessment of various risks, that is, for the evaluation of probabilities. It follows that the evaluation of probabilities is, in many cases, a step in the formation of preferences. Asking Carol what probabilities she assigns to various events by observing her own preferences is a circular proposition that is going to advance Carol neither in the evaluations of probabilities nor in the formation of preferences.

Observe also that Carol can't recall her past experiences with the two items on offer and compare her satisfaction with them, as Dan could do in choosing ice cream. Carol has not lived through a lifetime of work and retirement before. Even relying on other people's experience might not suffice, because no one has yet lived under the economic conditions that will prevail when Carol retires. Indeed, Carol may go through a complicated mental process, and possibly consult friends and colleagues, before choosing. And unlike Dan's ice cream decision, it is quite possible that someone could convince Carol that her tendency to prefer, say, the company pension plan was "wrong" by explaining to her the consequences of that choice if, say, the company were to be taken over by another firm. In a simple consumer problem, such as Dan's, preferences over different kinds of ice cream can be taken as "primitive". Primitive preferences need not be immutable: they might change in the long run as a result of advertisement, new information, or habit formation. Yet, such preferences exist, are available to introspection without the intervention of reasoning, and for many applications they also appear to be stable in the short run. By contrast, in Carol's example, preferences cannot be primitive. Rather, Carol's preferences between pension plans are the result of reasoning, or even explicit calculations that depend on the probability of various events.

We therefore hold that a rational decision maker's preferences may not satisfy Savage's axioms. The converse is also true: we may be unwilling to classify as "rational" every decision maker who does satisfy the axioms. The reason is that the axiomatic approach places constraints neither on the subjective probabilities, nor on the utility function that can be used to represent preferences. Such an agnostic position may be reasonable in the case of utility functions; indeed, most introductory textbooks emphasize that consumer

rationality is only a matter of internal consistency, and does not restrict the consumer's tastes in any fundamental way. But a similar agnostic position for the case of probability is far less defensible. Because probabilistic beliefs can be in accordance with evidence or at odds with it, some beliefs are more sensible than others. For example, assume that Ann's beliefs, as reflected in her choices, assign high probability of success to candidates graduating from a certain school, despite the fact that such candidates consistently fail in the program. Such beliefs could be considered "unreasonable". Similarly, many beliefs in supernatural phenomena and many superstitions are considered "irrational" because they conflict with evidence. Yet, nothing prevents a decision-maker from holding such beliefs and also satisfying Savage's axioms. In other words, Savage's axiomatic system restricts choices only to be internally coherent, and is therefore insufficient for an intuitive definition of rationality. Gilboa, Postlewaite, and Schmeidler (2006) develop this argument in some detail.

Despite the appeal of the axiomatic justification of subjective expected utility maximization, the foundation of subjective probabilities remains unsatisfactory for economic modeling. A decision-maker may seek guidance in Savage's axiomatic derivation to form probability beliefs, but the axioms are significantly less plausible than they seem at first glance. Specifically, the decision-makers involved may find that they have no a priori preferences over the relevant gambles, or that their preferences tend to violate the Sure Thing Principle.

This conclusion leads to two questions. First, are there alternatives to modeling decision-makers as maximizing expected utility with respect to given probability beliefs? Second, if a decision-maker's beliefs do not come from preferences over lotteries, what is the basis for beliefs? We address these questions in turn.

An Alternative Approach: The Multiple Prior Model

Models that assume that decision-makers maximize expected utility have been tremendously useful in generating insight into economic behavior under uncertainty. This fact is not contradicted by empirical or experimental violations of expected utility theory. Indeed, all economic models fail to be perfectly accurate descriptions of reality. When a useful model is shown to be inaccurate the key question becomes, is it inaccurate in an important way? Can it lead to qualitatively wrong conclusions? And, if so, what alternative models might provide better guidelines for understanding economic situations and generating predictions about them?

Consider again the example above in which Ann wants to assign a probability to a candidate from a foreign country being successful when she has no experience with students from the candidate's college or from his country. We suggested above that, out

of ignorance, Ann might assign a probability of success equal to the general success rate, say 60 percent. But Ann knows that this number is a default rate in which she has little confidence. She might feel uneasy about making decisions that hinge on this particular probability. In fact, she may even feel that it would not be rational to rely on this assessment, which is somewhat arbitrary. An alternative that Ann might consider is to model her state of knowledge more explicitly, and to give up on the notion that she should assign an exact probability to the event that this student will be successful. Instead, she can assign a range of probabilities to this event – say, between 55 and 65 percent.⁵ Ann might then calculate her expected utility for each of the probability distributions and make her decision based on the *set* of expected values she obtains. For example, if she wishes to be cautious, she might assign to each alternative the minimum expected value over her set of priors and choose that alternative that yields the highest such minimum. Gilboa and Schmeidler (1989) axiomatized the multiple prior model discussed here. Others models employing multiple priors include Bewley (2002), Klibanoff, Marinacci, and Mukerji (2005), Seo (2007), and Maccheroni, Marinacci, and Rustichini (2006a,b). These models are also axiomatically-based, that is, they have characterizations of their respective decision rules by the patterns of behavior that are compatible with them.

The multiple prior model in which the decision maker chooses the alternative that maximizes the minimum expected utility across the possible priors is clearly an alternative to the standard expected utility model. Moreover, ranges of probabilities can reflect the intuitive notion that a decision-maker might feel more confident in some beliefs than in others. But does the presumably more intuitive multiple prior model lead to new insights? The answer is “yes”.

As one example, the multiple prior model has been applied to the problem of optimal investment. Suppose that Carol has 100 shares of Intel stock, and, in addition, several thousand dollars in the bank. Carol must decide whether she should sell some of her Intel stock, buy more, or leave her portfolio as it is. In the standard expected utility model (setting aside considerations such as transaction costs or the possibility that the stock price itself offers information), there will be a unique price at which Carol would neither buy nor sell any stock; at any higher price she would sell some or all the stock, and at any lower price she would buy more. This conclusion seems overly sharp. It seems more likely that there would be a *range* of prices at which Carol would be willing to leave her portfolio unchanged. Dow and Werlang (1992) showed that in the multiple prior model of Gilboa and Schmeidler (1989) this is exactly what occurs. When Carol’s beliefs are represented as a *set* of priors rather than a single prior, she leaves her portfolio unchanged over a certain range of prices.

⁵ One could argue that Ann should have a “prior over the set of priors”, and simply reduce the problem to a single prior by computing compound probabilities. However, if Ann does not know what the prior probability should be, it seems even less likely that she would be confident about a specific prior over the set of possible priors.

To see the basic logic of this result, suppose that there are two states of the world, and an asset A that yields \$1 in state 1 and -\$1 in state 2. Consider an investor who holds none of the asset but can buy or sell short a positive quantity of the asset, and who believes that the probability of state 1 is in the range $0.4 \leq p \leq 0.6$. For simplicity, assume that the utility function is $u(x)=x$. If the investor evaluates a purchase of an additional unit of the asset A, the lowest expected utility is obtained at $p=0.4$, and it is -0.2. On the other hand, if the investor evaluates a sale of a unit of the asset, the lowest expected utility, obtained this time at $p=0.6$, yielding -0.2 as well. Thus, if the price of the asset is in the range $(-0.2, 0.2)$, the investor will prefer neither to buy nor to sell A.

The multiple prior model has also been employed in a job search model by Nishimura and Ozaki (2004). They ask how an unemployed agent will react to increasing uncertainty in the labor market. In a Bayesian model, greater uncertainty might be captured by higher variance of the job offers the agent receives. Other things being equal, an increase in variance (holding the mean constant) should make the agent less willing to accept a given offer, knowing that he has a chance to receive a substantially better offer in the future. However, the counter-intuitive conclusion that greater uncertainty should cause an agent to be more likely to turn down a current offer is the result of the assumption that all uncertainty is quantifiable by a probability measure. Nishimura and Ozaki show that in a multiple prior model – assuming, again, an “uncertainty averse” agent who uses the “maxmin” rule of choosing the option whose worst case expected utility is the highest – the conclusion might be reversed: in the presence of greater uncertainty, modeled as a larger set of possible priors, agents will be more willing to take an existing job offer rather than bet on waiting for better ones in the future.

Hansen and Sargent (2001, 2003, 2006) have applied the multiple prior model to macroeconomic questions starting from the viewpoint that, whatever probability model a policymaker might have, it cannot be known with certainty. Considering a set of priors around a given model, and asking how robust economic policy would be to variations in the underlying probability, they revisit and question classical results. Hansen, Sargent and Tallarini (1999) compare savings behavior under expected utility maximization with savings behavior of a “robust decision-maker” who behaves in accordance with the multiple prior model. They show that the behavior of a robust decision-maker puts the market price of risk much closer to empirical estimates than does the behavior of the classical expected utility maximizer.

In short, the multiple prior model can yield qualitatively different and more plausible results than a standard Bayesian model. The latter is a good model for many situations, where theoretical results remain qualitatively unchanged if analyzed with a multiple prior model. However, there are also many theoretical results that hinge on the existence of a prior belief defined by a single probability. When a result seems to depend on such a prior, one may become suspicious about its general applicability.

Belief Formation

Savage's axioms might convince Ann and Bob that they wish to maximize expected utility relative to a subjective prior, and the axioms may help Ann and Bob to elicit such a prior if they have complete preferences that satisfy the axioms. But decision-makers who attempt to define these preferences via the formation of their beliefs will find little help in the axioms. The question remains, therefore, how should Ann and Bob define prior beliefs? What are reasonable beliefs to hold?

Let us start with Ann. Assume that, trying to predict the success of a new candidate, she wants to use all the available data about successes and failures of Ph.D. students in the past. However, she would also like to give more similar candidates more weight in her probability assessment. A simple way to do this would be to choose a "similarity function", which would measure the degree to which two candidates are similar when evaluated for possible admission, and compute the relative frequency of success among students in the program, where each student is weighed by his or her similarity to the new candidate. Thus, the "probability" of success of a candidate would be the sum of similarities to all students who have succeeded, divided by the sum of her similarities to all students who either succeeded or failed.

If the similarity function is constant, the formula reduces to standard empirical frequencies, giving equal weight to all past observations, ignoring differences in the degrees of similarity between these past cases and the one in question. If, however, the similarity function takes the value 1 for cases that are identical to the new one in all measurable variables, and the value 0 otherwise, the similarity-weighted formula becomes the empirical frequency in the sub-database defined by the values of these variables. A similarity-weighted relative frequency thus suggests a continuous spectrum between these two extremes. This approach has been axiomatized in Billot, Gilboa, Samet, and Schmeidler (2005) and in Gilboa, Lieberman, and Schmeidler (2006), building on ideas of case-based decision theory (Gilboa and Schmeidler, 2001).

An obvious question about this approach is that it may appear that the problem of finding an appropriate probability has simply been replaced by the problem of finding an appropriate similarity function. However, reducing the question of probability to similarity is a meaningful step. In particular, an "objective" similarity function may be computed from existing data, by finding the function that best fits the data, if we were to use the similarity-weighted frequency as our prediction formula. We can illustrate this process with Ann's problem.

Ann would start by choosing a certain parameterized family of similarity functions. For each function in this family, Ann could imagine going over her database of Ph.D. students, and asking for each observation i , what her prediction would have been if she were asked to predict the outcome of case i , given all other cases in the database, and

using the specific similarity function at hand. She would then select the similarity function that minimizes the sum of squared errors of these hypothetical predictions relative to the observed realizations. The similarity function that minimizes the sum of squared errors would then be used to assess probability of success in the next observation. Gilboa, Lieberman, and Schmeidler (2006) refer to the resulting function as the “empirical similarity”.

This approach is offered as a model of the way in which the decision-maker forms beliefs. There are several benefits to such an approach. First, this approach emphasizes that beliefs do not typically arise from introspection. Rather, beliefs often result from conscious assessments by the decision-maker. Those assessments take as raw data past experiences, observations, and conversations with others. They may not always follow a precise calculation as offered by the model, but they do involve the weighing of different observations and the learning of the weights themselves from experience as well. Opening what is often taken to be a black box and analyzing the process of belief formation allows us to make inferences about *which* beliefs are more reasonable than others for a decision-maker. In particular, many economic problems have to do with equilibrium selection, as in the case of predicting market bubbles and crashes. The problem of equilibrium selection is fundamentally a problem of belief formation: the equilibrium that will be played is the equilibrium that is believed to be played. If we can predict which beliefs will be generated, we can also predict which equilibrium will result.

Second, all other things equal, models that are more descriptively accurate are preferable. The primary test of a model is the insight from the predictions of the model. But the way in which a problem is modeled also affects the subsequent questions that a researcher might ask. If instead of taking beliefs as exogenous, researchers also model the formation of those beliefs, they will be led to question the relative importance of different factors on beliefs. Shea (2004) asks whether the beliefs of children whose parents were financially devastated by the Great Depression systematically differ from those whose parents were spared? Do people who have been laid off put greater probability on there being a serious recession in the next five years? These are examples of questions of fundamental importance that require modeling the formation of beliefs.

Modeling the belief formation process may also help in understanding when a particular model of beliefs is applicable. Specifically, consider Bob’s problem. If we were to apply the similarity-weighted frequencies approach to this problem, the approach may seem less satisfactory than in Ann’s problem. For instance, the fact that the United States was involved in two Gulf Wars in the past two decades need not make it more plausible that the United States will be involved in another such war in the coming decade. In fact, the causal relationships between consecutive wars are intricate enough to support a variety of conclusions about the possibility of another such war. In the first problem, Ann had data about many students, whose successes and failures could be assumed causally independent from each other. Even though the students were not identical, certain notions of probability could be defined, be it by the empirical similarity weighted frequencies or by logistic regression. By contrast, Bob is facing a problem with relatively

few cases, which are anything but causally independent. There may be no rational, reasoned way to assign probabilities in Bob's problem, and one may have to do with less structured models (such as the multiple prior model mentioned above). The exercise of modeling belief formation does not help us only in finding reasonable beliefs to be used in our models; it can also help us delineate the scope of applicability of probabilistic models.

The Bayesian approach does not suggest a model of the formation of prior beliefs, but it describes the way that these beliefs are updated – namely, according to Bayes's rule. While our main point is that economists should delve into the prior-formation process, we also note that Bayesian updating is unlikely to be the only way in which beliefs are updated. In particular, Bayesian updating cannot account for adjustments of beliefs in the absence of new information. There are many instances in which people adjust their beliefs in the face of arguments that do not present new information, but suggest that different conclusions should be drawn from the same body of knowledge. Aragonès, Gilboa, Postlewaite, and Schmeidler (2005) argue that, due to computational complexity considerations, such “fact-free” learning is unavoidable even if agents are rational. This argument suggests that economists should be interested in realistic models of the way beliefs are formed, as well as of the way they are updated.

Conclusion

Models relying on the standard expected utility approach, which identifies a set of outcomes and places a probability on each outcome, have led to fundamental insights in virtually all fields of economics. At the same time, we find this approach lacking in two important ways, both of which may restrict the insights gained from formal models.

The first limitation is that the standard expected utility model, along with Bayesian extensions of that model, restricts attention to beliefs modeled by a single probability measure, even in cases where no rational way exists to derive such well-defined beliefs. We have argued that allowing more general models of beliefs may lead to more realistic results. The second limitation of the Bayesian framework is its silence regarding the origins of beliefs. We argued that a better understanding of the process by which beliefs are generated, whether these beliefs are Bayesian or not, may help researchers in providing better predictions, as in the case of equilibrium selection problems.

Both limitations of the standard expected utility approach and the Bayesian approach are partially dealt with in classical statistics. The classical statistical inference problem is defined by a *set* of probabilities, or distributions. Point and interval estimation consist in choosing a distribution or a set of distributions (respectively) from the original set. Hypothesis testing deals with the possibility of trimming the original set to a subset

thereof. Relatedly, statistics explicitly models the formation of beliefs based on data. Empirical work in economics typically employs classical (non-Bayesian) statistical techniques, such as hypothesis testing and interval estimation. As such, it implicitly adopts an approach in which a set of probability models is considered possible, and observations are used to trim it down. Economic theory may benefit from allowing economic agents to be non-Bayesian as well, and from explicitly modeling the way people form beliefs based on observations.

Acknowledgements

Our thinking on these issues was greatly influenced by discussions with many people. We are particularly grateful to the editors who have provided help and advice throughout the preparation of this paper.

Gilboa and Schmeidler gratefully acknowledge financial support from Israel Science Foundation Grants Nos. 975/03 and 355/06, the Pinhas Sapir Center for Development, and the Polarization and Conflict Project CIT-2-CT-2004-506084 funded by the European Commission-DG Research Sixth Framework Programme. Postlewaite gratefully acknowledges support from the National Science Foundation, Grants Nos. SES-05-27656 and SES- 03-50969-002

References

- Anscombe, Francis J. and Robert J. Aumann (1963) "A Definition of Subjective Probability", Annals of Mathematical Statistics, 34:199-205
- Aragonès, Enriqueta, Itzhak Gilboa, Andrew W. Postlewaite, and David Schmeidler, "Fact-Free Learning", American Economic Review, 95 (2005), 1355-1368.
- Bewley, Truman (2002), "Knightian Decision Theory: Part I", Decisions in Economics and Finance, 25: 79-110. (Working paper, 1986)
- Billot, Antoine, Itzhak Gilboa, Dov Samet, and David Schmeidler, "Probabilities as Similarity-Weighted Frequencies", Econometrica, 73 (2005), 1125-1136.
- de Finetti, Bruno (1937), "La Prevision: Ses Lois Logiques, Ses Sources Subjectives", Annales de l'Institute Henri Poincare, 7: 1-68.
- Dow, James, and Sergio R. C. Werlang, "Uncertainty aversion, risk aversion, and the optimal choice of portfolio", Econometrica, 60 (1992), 197-204.
- Drèze, Jacques (1961), "Les fondements logiques de l'utilité cardinale et de la probabilité subjective", in La Decision, Colloques Internationaux du CNRS, Paris, 73-97.
- Ellsberg, Daniel (1961), "Risk, Ambiguity and the Savage Axioms", Quarterly Journal of Economics, 75: 643-669.
- Gilboa, Itzhak, Andrew W. Postlewaite, and David Schmeidler (2006), "Rationality of Belief, or: Why Savage's Axioms are Neither Sufficient nor Necessary for Rationality", mimeo.
- Gilboa, Itzhak, and David Schmeidler (1989), "Maxmin Expected Utility with a Non-Unique Prior", Journal of Mathematical Economics, 18, 141-153. (Working paper, 1986)
- Gilboa, Itzhak, and David Schmeidler (2001), A Theory of Case-based Decisions, Cambridge University Press, 214 pages.
- Gilboa, Itzhak, Offer Lieberman, and David Schmeidler, "Empirical Similarity", Review of Economics and Statistics, 88 (2006), 433-444.
- Hacking, Ian (1975), The Emergence of Probability. Cambridge: Cambridge University Press.

- Hansen, Lars Peter, and Thomas J. Sargent (2001), “Acknowledging Misspecification in Macroeconomic Theory”, Review of Economic Dynamics, 4: 519–535.
- _____ (2003), “Robust Control of Forward-Looking Models”, Journal of Monetary Economics, 50: 581–604.
- _____ (2006), Robustness. Princeton: Princeton University Press, forthcoming.
- Hansen, Lars Peter, Thomas J. Sargent and Thomas D. Tallarini (1999), “Robust Permanent Income and Pricing”, Review of Economic Studies, 66: 873–907.
- Karni, Edi (1993), “A Definition of Subjective Probabilities with State–Dependent Preferences”, Econometrica, 61 187–198.
- _____ (1996), “Probabilities and Beliefs”, Journal of Risk and Uncertainty, 13: 249–262.
- _____ (1999), “Elicitation of Subjective Probabilities when Preferences are State–Dependent”, International Economic Review, 40: 479–486.
- _____ (2003), “On the Representation of Beliefs by Probabilities”, Journal of Risk and Uncertainty, 26:17–38.
- Karni, Edi, David Schmeidler, and Karl Vind (1983), “On State Dependent Preferences and Subjective Probabilities”, Econometrica, 51: 1021–1032.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2005), “A Smooth Model of Decision Making under Ambiguity”, Econometrica, 73: 1849–1892.
- Knight, Frank H. (1921), Risk, Uncertainty, and Profit. Boston, New York: Houghton Mifflin.
- Lichtenstein, Sarah and Paul Slovic (1971), “Reversals of Preferences between Bids and Choices Gambling Decisions”, Journal of Experimental Psychology, 89: 46–55.
- Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini (2006a), “Ambiguity Aversion, Robustness, and the Variational Representation of Preferences”, Econometrica, 74: 1447–1498.
- _____ (2006b), “Dynamic Variational Preference”, Journal of Economic Theory, 128: 4–44.
- Nishimura, Kiyohiko and Hiroyuki Ozaki (2004), “Search and Knightian Uncertainty”, Journal of Economic Theory, 119: 299–333.
- Ramsey, Frank P. (1931), “Truth and Probability”, in The Foundation of Mathematics and Other Logical Essays. New York, Harcourt, Brace and Co.

- Savage, Leonard J. (1954), The Foundations of Statistics. New York: John Wiley and Sons. (Second ed., 1972, Dover)
- Schmeidler, David (1989), “Subjective Probability and Expected Utility without Additivity”, Econometrica, 57: 571-587. (Working paper, 1982)
- Seo, Kyoungwon (2007), “Ambiguity and Second-Order Belief”, mimeo.
- Shafer, Glenn (1986), “Savage Revisited”, Statistical Science, 1: 463-486.
- (1978), “Non-additive probabilities in the work of Bernoulli and Lambert”, Archive for History of Exact Sciences, 19: 309-370.
- Shea, John (2003), “Childhood Deprivation and Adult Wealth”, mimeo.
- Tversky, Amos, & Kahneman, Daniel (1974), “Judgment under uncertainty: Heuristics and biases”, Science, 185: 1124-1131.
- von Neumann, John, and Oskar Morgenstern (1944), Games and Economic Behavior. New York: Wiley.