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 07-023

"Probabilities in Economic Modeling"

by

Itzhak Gilboa, Andrew Postlewaite and David Schmeidler

http://ssrn.com/abstract=1004727

Probabilities in Economic Modeling*

by

Itzhak Gilboa[1], Andrew Postlewaite[2], and David Schmeidler[3]

August 2007

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.

1. Where do beliefs come from?

The standard practice in economics when modeling the behavior of economic agents facing uncertainty is to assume that the agents are Bayesian. That is, people are assumed to have probabilistic beliefs over any source of uncertainty, to update these beliefs in accordance with Bayes's rule, and to use them in decision making, typically as a basis for expected utility maximization.

* Our thinking on these issues was greatly influenced by discussions with many people. 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.

[1] Tel-Aviv University, HEC, and Cowles Foundation, Yale University. igilboa@tau.ac.il

[2] University of Pennsylvania. apostlew@econ.upenn.edu

[3] Tel-Aviv University and Ohio State University. schmeid@tau.ac.il

The Bayesian paradigm is an elegant and coherent way to deal with uncertainty. Yet, it is not always clear how should probabilistic beliefs be formed. The following 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 US 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". Unfortunately, there are no agreed-upon methods for assigning probabilities to the events that concern them. Laplace's "principle of indifference"[4] suggesting equal probabilities to all possible eventualities, is obviously inapplicable in these problems. Empirical frequencies are also 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. On the other hand, considering only students who are identical in all observable parameters to the candidate at hand would likely leave Ann with an empty sample.

Likewise, Bob might attempt to employ an empirical frequency approach for his problem. He could 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, and other 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.

Thus, in both examples, it is not obvious how probabilities should be defined. Moreover, it appears that this will be the case in many economic problems of interest. Probabilities are actually "given" in only very restricted situations such as state lotteries or casino games. There are many situations such as insurance problems where probabilities are not stated in the problem, but where they can be reasonably approximated by relative frequencies of comparable instances computed from publicly available data. But in a vast

[4] This is also known as the "principle of insufficient reason".

range of economic problems probabilities are neither explicitly given nor can they be approximated by relative frequencies. 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.

The remainder of this paper is organized as follows. Section 2 considers the notion of subjective probabilities, which is the standard approach to the definition of probabilities in examples such as the above. We discuss this approach and its limitations. We argue that rational decision makers may violate the axioms underlying the Bayesian approach, and that the axioms are, in this sense, too restrictive. Conversely, we also argue that the axioms are too general because they do not restrict probabilistic beliefs in any way.[5] The following two sections are devoted to the implications of these limitations to economic modeling. Section 3 deals with alternative models of decision making, which may account for a wider set of phenomena than the Bayesian one. We discuss a few examples of economic models that make use of such approaches. Section 4 calls for the development of formal, explicit theories of the belief formation process, and attempts to identify the scope of applications in which the Bayesian approach might be useful. To conclude, our claim is that the Bayesian approach is too restrictive in some applications, and too general in others, and that a theory of belief formation may help both to refine the predictions of economic models when agents are Bayesian, and to improve these predictions when they are not.

2. Subjective probabilities

The Bayesian paradigm holds that individuals should have subjective probabilities even when objective probabilities cannot be defined. The paradigm relies on a seemingly compelling axiomatization proposed by Savage (1954).[6] This axiomatization considers observable choices between pairs of uncertain acts and rules out patterns of choices that do not seem "sensible". To understand the 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 simple

[5] 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. In this note we deal with individual (as opposed to interactive) decision making.

[6] There are other axiomatizations of subjective probability, coupled with the principle of expected utility maximization (see, e.g., Ramsey (1931), de Finetti (1937), and Anscombe-Aumann (1963)). We focus on Savage's axiomatization because it is widely perceived to be the most satisfactory from a conceptual viewpoint.

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 her choices by maximization of her subjective expected utility.[7]

Savage postulated four conceptually important axioms.[8] 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 can choose between them.[9] *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, but are not the main focus of this paper and we therefore do not discuss them here.[10]

A key 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 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 she prefers G3 to G4.

---

[7] The term "subjective" in "subjective expected utility" refers to the probability measure. Of course, the utility function is also subjective, as in classical consumer theory, as well as in von-Neumann and Morgenstern's (vNM, 1944) axiomatization of expected utility maximization. However, vNM assume that probabilities are given, and hence their theory is restricted to such situations. Savage's main contribution was to extend the expected utility paradigm to situations where objective probabilities may not exist.

[8] Savage also needed three additional axioms, which can be viewed as "technical", guaranteeing notions of continuity and ruling out trivial cases. These axioms will not be discussed here.

[9] The decision maker can be indifferent between two bets.

[10] See Drèze, (1961), Karni, Schmeidler, and Vind, (1983), and Karni (1993, 1996, 1999, 2003).

These, and the other Savage 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 she 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 she maximized the sum of the utilities of the outcomes weighted by the probabilities she will get them.

Savage's theorem that a person behaves as though she 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 also 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 testing subjective expected utility theory. Because direct empirical tests of the theory may be fraught with identification problems, 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 rather 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 that she should accept expected utility theory as a descriptive theory of the decisions that will be made to the same degree that she 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 simpler, 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 comparison will identify a unique p that can be defined as her subjective probability for the event in question.

Whether interpreted normatively or descriptively, Savage's axioms are viewed as standards of rationality. The Bayesian approach therefore holds that, for Ann and Bob to be rational, there must be probabilities that represent their choices via the expected utility formula. Furthermore, if these decision makers do not know their subjective probabilities, Savage's axiomatization suggests 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% 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 for a moment that this general success rate is also 60%. Yet, Ann knows that this number, 60%, was assigned almost as a default. By contrast, the 60% 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 than with one whose distribution is not known.

Since the early days of probability theory there has been a distinction between "chances", referring to known probabilities as in parlor games, and "probabilities", used to describe subjective degrees of belief.[11] The Bayesian approach, logically necessitated by Savage's axioms, fails to distinguish between probabilities that are based on data and probabilities that result from 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 the two distributions feel different, and, as a result, our willingness to bet on them need not be the same.[12] In a classic experiment, Ellsberg (1961) has shown that people often express preferences for bets with known probabilities over bets with unknown probabilities.[13]

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.[14] 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.

But in the presence of uncertainty, whether about objective outcomes or about one's subjective experience, completeness of preferences is less compelling. Assume that Carol is taking a new job, and she is offered either the company pension plan or a 409B plan. Carol must make a choice, as does Dan who has to choose what flavor of ice cream to have. But Carol's choice is very different from Dan's. If Carol attempts to introspect, she is likely to find out that she has no a-priori preferences between the plans, and she has to think about them in order to make a decision. This is a case in which reasoning has to

---

[11] See, e.g., Hacking (1975) and Knight (1921).

[12] This example is from Schmeidler (1989), which provides a discussion of the issue.

[13] There are other reasons why Savage's 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 (see Drèze, 1961, Karni, Schmeidler, and Vind, 1983, Karni, 1993, 1996, 1999, 2003).

[14] See Shafer, 1986, Bewley, 2002, and Gilboa, Postlewaite, and Schmeidler, 2006.

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 extend to complicated choices between uncertain prospects. In particular, the two pension plans would have to be *analyzed* for 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 simply recall her past experiences with the two items on offer and compare her satisfaction with them as Dan would in choosing ice cream. Carol has not lived through retirement before, in a way that would allow her direct experience with the outcomes to guide her choice between actions. 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. In short, "rational" decision making in this example requires reasoning to generate preferences. Indeed, Carol would probably 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 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.[15] By contrast, in Carol's example, preferences cannot be primitive. Rather, they are the result of reasoning, or even explicit calculations that depend on the probability of various events.

These descriptive problems with the use of expected utility theory do not imply that it is useless. On the contrary, the value of a normative theory is revealed precisely when it fails descriptively. Thus, if Ann or Bob were to violate the Sure Thing Principle and/or the completeness axiom, one may ask them to define their preferences in such a way that they satisfy Savage's axioms. As pointed out above, one way to guarantee that the axioms hold is to define preferences by expected utility maximization. In fact, this is the only known algorithm that guarantees this result. Ann and Bob should therefore define their utility functions and their subjective probability measures. It is important to observe that the axiomatic derivations of expected utility theory constrain neither the subjective probabilities one should use, nor one's utility function. *Any* choice of a utility function and a probability measure define, via expected utility maximization, a preference relation that satisfies the axioms. The axioms therefore are of no help to a decision maker who wishes to know what are *reasonable* probabilities to use in a given decision problem.

---

[15] Such preferences might still 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.

The silence of the axioms regarding the choice of utility and probability might indicate that the axioms are also unsatisfactory for descriptive purposes. The axiomatizations of subjective expected utility maximization might appear to be analogous to those of utility maximization in consumer theory: the latter imply the existence of a utility function whose maximization is observationally equivalent to the consumer's choices, but they say little about the specific utility functions that consumers might (or should) have. Such an agnostic position may be reasonable in the case of utility functions; indeed, most introductory textbooks emphasize consumer sovereignty. But a similar agnostic position for the case of probability is far less defensible. Because 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 of their conflict with evidence. Yet, there is nothing that prevents a decision maker from holding such beliefs and satisfying Savage's axioms. Savage's axiomatic system discussed above, which restricts choices only to be internally coherent, is therefore insufficient for an intuitive definition of rationality.[16]

To conclude, despite the appeal of the axiomatic justification of subjective expected utility maximization, the basis of subjective expected utility maximization 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 immediate 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 her preferences over lotteries, what is their basis? We address these questions in turn.

3. Generalizing the Bayesian model

Models that assume that decision makers maximize expected utility have been tremendously useful in generating insight into economic behavior under uncertainty. The fact that the expected utility model is not perfectly accurate is neither surprising nor necessarily relevant. Indeed, all economic models fail to be perfectly accurate descriptions of reality. The question is, are they inaccurate in an important way? Can they lead us to qualitatively wrong conclusions? And, if so, what alternative models might provide better guidelines for understanding economic situations and generating predictions about them?

---

[16] Gilboa, Postlewaite, and Schmeidler (2006) develops this argument in some detail.

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, she might assign a probability of success equal to the general success rate, say 60%. But Ann knows that this number is a default rate that she has little confidence in. 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 explicitly model her state of knowledge, 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%.[17] 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.[18]

This is clearly an alternative to the standard expected utility model. Moreover, the range of probabilities used can reflect the intuitive notion that a decision maker might feel more confident in some beliefs than in others. But the question is, does it matter? Does the presumably more intuitive multiple prior model lead to new insights?

The answer to this question is "Yes". For 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 exactly as it is. In the standard expected utility model, 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.[19] This conclusion seems counter-intuitive. 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, there will be a range of prices at which she leaves her portfolio unchanged.

[17] 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.

[18] Gilboa and Schmeidler (1989) analyzes the multiple prior model discussed here. Others have analyzed other multiple prior models, e.g., Bewley (2002), Klibanoff, Marinacci, and Mukerji (2005). Most of these models are axiomatically-based, that is, they have characterizations of their respective decision rules by the patterns of behavior that are compatible with them.

[19] We ignore here other considerations such as transaction costs or the possibility that Ann might learn something from the price, and the information was just enough to make her want exactly the amount of stock she currently holds.

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 get better ones later on. This counter-intuitive conclusion is a result of the assumption that all uncertainty is quantifiable by a probability measure. Nishimura and Ozaki (2004) show that in a multiple prior model (assuming, again, an "uncertainty averse" agent who uses the maxmin rule) 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 is the 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.

The preceding examples illustrate how the multiple prior model can yield qualitatively different and more plausible results than a Bayesian model. Importantly, when a theoretical result seems to hinge on the existence of a single prior, one may become suspicious about its general applicability.

4. Belief formation

We argued above that the Savage framework doesn't guide Ann and Bob in determining their beliefs. Savage's axioms might convince Ann and Bob that they wish to maximize expected utility relative to a subjective prior, and they may help them 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 PhD 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, measuring 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 similarity to the new candidate. Thus, the "probability" of success of a candidate would be the sum of her similarities to all students who have succeeded, divided by the sum of her similarities to all students who either succeeded or failed.

This similarity-weighted formula generalizes both standard and conditional empirical frequencies. 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. The similarity-weighted relative frequency thus suggests a continuous spectrum between these two extremes.

Thus, we offer a modified frequentist approach as a possible solution to the belief formation question. This modified frequentist approach presumes that the decision maker has a similarity function that provides the weights she should put on past experiences when assessing a new event. But where does the similarity function come from? It would 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 PhD 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.[20]

There are several benefits to modeling the process by which a decision maker forms her beliefs. First, it emphasizes that beliefs do not typically arise from introspection; they most likely are the result of the decision maker's conscious calculation. Those calculations take as raw data past experiences, observations, and conversations with others. 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, there are many economic problems having to

[20] Gilboa and Schmeidler (2001) propose a "case-based" decision making model based on such similarity functions as an alternative to the Bayesian decision model.

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 be player.

Second, all other things equal, we prefer models that are more descriptively accurate. The primary test of a model is the insight we get from the predictions of the model. But the way we model a problem also affects the subsequent questions we ask. If instead of taking beliefs to be exogenously given and outside the purview of the analyst, we model the formation of those beliefs, we are immediately led to question the relative importance of different factors on beliefs. Do the beliefs of children whose parents were financially devastated by the great depression systematically differ from those whose parents were spared?[21] Do people who have been laid off put greater probability on there being a serious recession in the next five years? The answers to these questions are of fundamental importance and require seriously modeling the formation of beliefs.

Modeling the belief formation process may help us understand when a particular model of beliefs is applicable. Specifically, consider Bob's problem. If we were to apply the same similarity-weighted frequencies approach to this problem, we may find that it is less satisfactory than in Ann's problem. For instance, the fact that the US was involved in two Gulf Wars in the past two decades need not make it more plausible that the US 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, some notion of probability could be defined. By contrast, Bob is facing a problem with relatively few cases, which are anything but causally independent. It may well be the case that there is 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 on 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. Aragones,

[21] See Shea (2004) for a treatment of this question.

Gilboa, Postlewaite, and Schmeidler (2005) argue that, due to computational complexity considerations, such “fact-free” learning is unavoidable even if agents are rational. This arguments suggests that economists should be interested in realistic models of the way beliefs are formed, as well as of the way they are updated.

5. Conclusion

The Bayesian approach to modeling decision making under uncertainty has been tremendously fruitful in economics. Models relying on the approach have led to fundamental insights in virtually all fields of economics. At the same time, we find the Bayesian approach lacking in two important ways, both of which may restrict the insights gained from formal models.

The first limitation of the Bayesian framework is that it restricts attention to beliefs modeled by a single probability measure, even in cases where there is no rational way to derive such well-defined beliefs. As mentioned above, 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 us in providing better predictions, as in the case of equilibrium selection problems.

Both limitations of the Bayesian approach are partially dealt with in classical statistics. The classical statistical inference problem is defined by a *set* of probabilities, from which a subset (or a singleton) is to be chosen. The classical approach does not require the selection of a single probability, nor a probability over probabilities. And classical statistics explicitly models the formation of beliefs based on data. Empirical studies in economics, as in any other field of science, tend to adopt a classical rather than a Bayesian point of view. Economic theory can benefit from allowing economic agents to be non-Bayesian, and from explicitly modeling the way people form beliefs based on observations.

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)

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 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.

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 5l: l02l–l032.

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.

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)

Shafer, Glenn (1986), "Savage Revisited", Statistical Science, 1: 463-486.

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.