

*Rationality of belief or: why savage's
axioms are neither necessary nor sufficient
for rationality*

**Itzhak Gilboa, Andrew Postlewaite &
David Schmeidler**

Synthese

An International Journal for
Epistemology, Methodology and
Philosophy of Science


ISSN 0039-7857

Volume 187

Number 1

Synthese (2012) 187:11-31

DOI 10.1007/s11229-011-0034-2



Your article is protected by copyright and all rights are held exclusively by Springer Science+Business Media B.V.. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your work, please use the accepted author's version for posting to your own website or your institution's repository. You may further deposit the accepted author's version on a funder's repository at a funder's request, provided it is not made publicly available until 12 months after publication.

Rationality of belief or: why savage's axioms are neither necessary nor sufficient for rationality

Itzhak Gilboa · Andrew Postlewaite ·
David Schmeidler

Received: 30 September 2011 / Accepted: 30 October 2011 / Published online: 19 October 2011
© Springer Science+Business Media B.V. 2011

Abstract Economic theory reduces the concept of rationality to internal consistency. As far as beliefs are concerned, rationality is equated with having a prior belief over a “Grand State Space”, describing all possible sources of uncertainties. We argue that this notion is too weak in some senses and too strong in others. It is too weak because it does not distinguish between rational and irrational beliefs. Relatedly, the Bayesian approach, when applied to the Grand State Space, is inherently incapable of describing the formation of prior beliefs. On the other hand, this notion of rationality is too strong because there are many situations in which there is not sufficient information for an individual to generate a Bayesian prior. It follows that the Bayesian approach is neither sufficient nor necessary for the rationality of beliefs.

Keywords Probability · Rationality · Beliefs · Bayesian paradigm · Savage axioms

I. Gilboa (✉) · D. Schmeidler
Tel-Aviv University, Tel-Aviv, Israel
e-mail: tzachgilboa@gmail.com

I. Gilboa
HEC, Paris, France

D. Schmeidler
The Interdisciplinary Center (IDC), Herzliya, Israel
e-mail: davidschmeidler@gmail.com

D. Schmeidler
Ohio State University, Columbus, OH, USA

A. Postlewaite
University of Pennsylvania, Philadelphia, PA, USA
e-mail: apostlew@econ.upenn.edu

1 Rationality of belief and belief formation

Economic theory is both the birthplace and the prime application of the rational choice paradigm. Throughout the 20th century, economics has relied on rationality, and refined the definition of rational choice, offering concepts such as subjective expected utility maximization and Nash equilibrium, which have proved useful in several other disciplines.

The paradigm of rational choice offers a notion of rationality that is much more flexible than the more traditional notions discussed in preceding centuries. In the past, many writers expected rationality to settle questions of faith, advance science, promote humanistic ideas, and bring peace on Earth. Such writers did not shy away from arguing what is rational and what is not, and to take a stance regarding what Rational Man should do, believe, and aspire to.

By contrast, economic theory in the 20th century took a much more modest and relativist approach to rationality. No longer was there a pretense to know what Rational Man should think or do. Rather, rationality was reduced to various concepts of internal consistency. For example, rational choice under certainty became synonymous with constrained maximization of a so-called utility function. By and large, economic theory does not attempt to judge which utility functions make sense, reflect worthy goals, or lead to beneficial outcomes. In essence, any utility function would suffice for an agent to be dubbed rational. More precisely, utility functions might be required to satisfy some mathematical properties such as continuity, monotonicity, or quasi-concavity, but these do not impose any substantive constraints on the subjective tastes of the economic agents involved. The mathematical properties of the utility function cannot imply, for instance, that it is irrational not to save for one's retirement. Defined solely on the abstract mathematical structure, these properties may be viewed as restricting the form of preferences, but not their content.

This minimalist requirement has two main justifications. The first is the desire to avoid murky and potentially endless discussions regarding the “true” nature of rationality. The second is that such a weak requirement does not exclude from the economic discussion more modes of behavior than are absolutely necessary. As a result, the theory is sufficiently general to be applied in a variety of contexts, for a wide variety of utility functions.

In this article we do not take issue with the notion of utility maximization or with the rationality of expected utility maximization given probabilistic beliefs over a state space. Our focus is on the existence of such beliefs and their justification. Our point of departure is that, when addressing the question of rationality of belief, economic theory adopted the same modest approach it has employed for the question of rationality of tastes.

In an attempt to avoid the question of what it is rational to believe, as well as not to rule out possibly strange beliefs, economic theory and decision theory have adopted a definition of rational beliefs that, like the definition of rational preferences, is based on internal consistency alone. This highly subjective view dates back at least to [de Finetti \(1937\)](#). In modern economic thought, a decision maker who satisfies [Savage \(1954\)](#) axioms, and behaves as if they entertain a prior probability over a state space, will be considered a rational decision maker under uncertainty, and may be viewed as having rational beliefs.

Such a relativist notion of rationality of belief is not the standard notion of “Bayesianism” in philosophy or in statistics.¹ Moreover, even for economic theorists it is hardly intuitive. Consider a graduate student who believes that he is among the best economists in the world. Assume that he assigns probability 1 to this event, and that he takes decisions so as to maximize his expected utility with respect to these views. In the face of new evidence (failing prelims for example), he employs Bayes’s rule to update his probability. But since he ascribes zero probability to the event that he is not a supremely gifted economist, his updated beliefs are that his professors are simply not sufficiently smart to recognize the depth and importance of his ideas. Throughout the process, the student may well satisfy all of Savage’s axioms, as well as the implicit axiom of Bayesian updating.² Yet, we may agree that the student needs to be treated as delusional. Indeed, in everyday parlance we make distinctions between rational and irrational beliefs, but decision theory is silent on this issue.³

Reducing rationality of utility to internal consistency may result from the ancient recognition that tastes are inherently subjective (*de gustibus non est disputandum*). But one may say more about beliefs than about tastes. Rationality does not constrain one to like or to dislike the smell of tobacco but rationality does preclude the belief that smoking has no negative health effects. Similarly, a person who buys a painting may be wrong about the probability that it will be stolen, but not about which painting pleases her more.

Defining rationality of beliefs by internal consistency alone allows the theory to apply to a wide array of beliefs. But this generality might be costly. First, when we refuse to address the question of which beliefs are rational, we may not notice certain regularities in the beliefs entertained by economic agents. Second, by restricting attention to the coherence of beliefs, one evades the question of the generation of beliefs. Indeed, economic theory offers no account of the belief formation process. Beliefs are supposedly derived from observed behavior, but there is no description of how the beliefs that generated the observed behavior arose in the first place.

Economic theory would benefit from a theory of belief formation, and, relatedly, from a classification of beliefs according to their rationality. A theory of belief formation could suggest a systematic way of predicting which beliefs agents might hold in various environments. Rationality of beliefs may serve as a tool for comparing the relative plausibility of different equilibria in economic models.⁴ Moreover, if we had a theory of how beliefs are formed, we might also use it to delineate the scope of competing models for representation of beliefs. In particular, we would be able to tell when economic agents are likely to entertain probabilistic beliefs, and when their beliefs should be modeled in other, perhaps less structured ways.

¹ See Carnap (1952); Lindley (1965); Jeffrey (2004).

² See Ghirardato (2002).

³ The notion of equilibrium in economics and in game theory may be viewed as an implicit definition of rational beliefs. That is, rational beliefs are those that coincide with equilibrium behavior. However, such a definition does not enlighten us about the process by which rational beliefs come into being.

⁴ In a sense, Cho and Kreps (1987) “intuitive criterion” and related belief-based refinements are among the few exceptions in which economic theory does dare to rule out certain beliefs on the grounds of irrationality. For a discussion of this literature, see Mailath et al. (1993).

Because the Bayesian approach ignores the belief formation process, it appears to be too weak a standard of rationality. The bulk of this article is devoted to the converse claim, namely, that the Bayesian approach may also be too demanding in certain circumstances, and that rational decision makers need not always be Bayesian. Thus, we hold that Bayesian beliefs are neither necessary nor sufficient for rationality of belief.

We do not define the notions of “belief” or of “rationality” in this article. Our main point is that “Bayesianism” does not satisfactorily capture the intuitive notion of “rational belief”. Specifically, we claim that there are situations where Bayesian beliefs will not appear rational in an intuitive sense, and, conversely, there are situations in which beliefs that seem intuitively rational will fail to be Bayesian.

We devote the next section to a definition of the concept of “Bayesianism” discussed and criticized in this article. Its main goal is to bridge a culture gap and explain to non-economists how economic theorists have come to think of “Bayesianism”. Section 3 argues that a rational agent who consciously evaluates the plausibility of events need not be Bayesian in this sense. Thus, Sect. 3 attempts to show that, on cognitive grounds, rationality need not imply Bayesianism. Section 4 is devoted to the claim that rational behavior implicitly defines Bayesian beliefs even if the agent is not aware of these beliefs. It argues that rational behavior under uncertainty also does not imply the existence of (implicit) Bayesian beliefs. Finally, we conclude by considering alternative approaches to the generation and representation of beliefs.

Two caveats are in order. First, some of the arguments mentioned in this article have appeared previously in the literature (see references below). Second, we discuss the notion of Bayesianism that is commonly used in economic theory. As we explain below, it is a much more extreme variant of the Bayesian faith than those found in statistics, philosophy, or computer science. Our critique is only directed at the economic version of Bayesianism.

2 What is Bayesianism?

When economic theorists think of the term “Bayesian”, they typically refer to four basic tenets:

Grand state space In Savage’s words, a state should “resolve all uncertainty”, which implies that all possibly relevant causal relationships, and all that is known about the way information is obtained, are also specified in a “state”. Everyday theoretical work often ignores this principle, thereby simplifying the definition of the state space. But whenever a conceptual difficulty is encountered, one resorts to a sufficiently elaborate state space to resolve it.

Prior probability (i) Whenever a fact is not known, one should have probabilistic beliefs about it. (ii) These beliefs should be given by a single probability measure defined over a state space in which every state resolves all relevant uncertainty.

Bayesian updating In light of new information, the Bayesian prior should be updated to a posterior according to Bayes’s law.

Expected utility When facing a decision problem, one should maximize expected utility with respect to one’s Bayesian beliefs (incorporating all information that one has gathered).

In statistics, philosophy, computer science, artificial intelligence, and related fields, Bayesianism typically means only the second and third tenets. The two other tenets call for explication.

We begin with the fourth tenet. The other disciplines mentioned earlier often discuss beliefs, evidence, induction, and learning without an explicit reference to decision making, and thus they need not always specify a decision theory that accompanies the representation of beliefs. As a result, they can do without the fourth tenet or any alternative thereof. By contrast, economics is inherently interested in decision making, and it uses the representation of beliefs in order to predict or explain behavior. The fourth tenet states how beliefs affect, or are at least reflected in behavior.

Conversely, these behavioral implications can be the basis of an axiomatization of the Bayesian approach, coupled with expected utility maximization, as in [Ramsey \(1931\)](#); [de Finetti \(1937\)](#), and [Savage \(1954\)](#). These works showed that if a decision maker satisfied certain assumptions on the ranking of alternatives in an uncertain environment, her behavior was consistent with the second and fourth tenets of Bayesianism above, namely, that she made decisions so as to maximize expected utility with respect to a probability measure. One may, however, also provide behavioral axiomatizations of the Bayesian approach with other decision theories, such as prospect theory ([Kahneman and Tversky 1979](#)).⁵ Both types of axiomatization could also be used, in principle, for the elicitation of subjective probabilities based on behavior data.

It is perhaps less natural that economics is the only discipline in which the dominant interpretation of “Bayesianism” adopts also the first tenet, according to which the state space is elaborate enough to be able to describe anything of relevance. The roots of this belief may go back to the writings of [de Finetti \(1937\)](#), who was famous for his religious zeal when it came to the Bayesian faith. [Savage \(1954\)](#) was referring to this interpretation of the Bayesian model, but seemed to be more pragmatic and skeptical about it.⁶ It seems that [Harsanyi \(1967, 1968\)](#) greatly contributed to the acceptance of the first tenet in economics. In these path-breaking works Harsanyi considered games of incomplete information, that is, games in which the utility functions and beliefs of the players may not be known to other players. He suggested to analyze them as regular games in which, at the first stage, nature moves and selects the “types”, namely, utilities and beliefs, of all players. In the following stages, players are called upon to make moves, where their information need not be perfect, as is standard in games. Thus, in a trivial but brilliant modeling strategy, Harsanyi showed that game theorists who could deal with an element of risk in the models could also deal with unknown utilities or beliefs. There was no conceptual difference between, say, not knowing whether the opponent in a game is altruistic, and not knowing which hand she was dealt in a game of Poker.

⁵ [Machina and Schmeidler \(1992, 1995\)](#) provide behavioral axioms that characterize a decision maker who can be ascribed a Bayesian prior, but whose decisions are guided by a functional that may be non-linear in probabilities. [Rostek \(2010\)](#) provides such an axiomatization for decision makers who maximize quantiles, such as the median.

⁶ In a tone that is half-caricaturizing his own argument, Savage explains that an “act” should be viewed as a strategy one adopts once and for all, before one is born. His discussion elsewhere (including the notion of “small worlds”) shows that he might not have approved of the extreme interpretation which is today prevalent in economic theory.

Harsanyi's idea could be described as going back in time before the players got their own identity, and imagining what kind of reasoning they might be engaged in at this "original position". This mental exercise can't fail to remind us of Rawls's (1973) "veil of ignorance". In fact, Harsanyi has already used the same modeling strategy in his defense of utilitarianism (Harsanyi 1953), where he related utilitarianism to expected utility maximization in the "original position".

Harsanyi's works have opened the way for game theory to deal with a host of economic phenomena that involved information asymmetries. Equipped with this tool, economics could deal with many phenomena that were not formally modeled in the past, to analyze various "market failures", and also to export game theoretic models to other disciplines. The belief that any interactive situation should be modeled as a "Bayesian game" has become standard in economics. Coupled with the assumptions that all players also share the same prior distribution over the state space, it has come to be known as the "Harsanyi Doctrine".

Since the Harsanyi Doctrine came to dominate economics, generations of economic theorists were brought up on the proposition that the four tenets of Bayesianism are the basic way to analyze any problem of social interaction under uncertainty, as well as any problem involving uncertainty in general. The Bayesian way of thinking offered a convenient framework that is almost algorithmic: one starts out with the state space, whose description should include anything that can possibly matter. In case of doubt about the relevance of any conceivable course of uncertainty, it is brought forth and modeled explicitly as part of the definition of a "state". Next, a prior is formed, which allows one to weed out irrelevant states by assigning them probability zero. Then the prior is updated based on new evidence, and if a decision has to be taken, it is taken so as to maximize expected utility.

The current state of the art is that most economic theorists accept this as the way to analyze any problem. For example, if we were to discuss the question of induction, and the formation of beliefs in rules from observations, the textbook economic approach would be to say that a state of the world specifies all observations that can ever be gathered, and thereby also the degree of accuracy of every rule than can ever be imagined. Hence, defining the states "correctly", there is no problem of induction—inductive inference is no more than Bayesian updating. This, at least, is how economists model reasoning by rational agents.

To consider another example, a well-trained economic theorist would not see the need for a concept such as Carnap's confirmation function (Carnap 1952), mapping evidence to probabilistic beliefs. Rather, she would adopt the Grand State Space, in which every state specifies what this function is, and model the reasoner's choice of the confirmation function by a prior belief over this state space. Similarly, whereas Bayesian statisticians would use a prior belief over a parameter, but base the choice of a prior on theoretical considerations as well as on past experience, economic theory would model all such past experience as being reflected in the definition of the states, and any theoretical considerations in the prior beliefs. If asked to justify this notion of rationality, the theorist would mention the axiomatic foundations—mostly Savage (1954)—as a convincing argument that, even when applied to an all-encompassing state space, agents should be Bayesian.

3 Insufficient information

All four tenets of Bayesianism have come under attack, especially as descriptive theories. The most common critique of Bayesian principles in economics is the descriptive failure of the last tenet, namely that people maximize expected utility when they are equipped with probabilistic beliefs. This critique began with Allais's famous example in which a large proportion of people make choices that are inconsistent with expected utility theory⁷, and continued with the behavioral deviations from EUT (expected utility theory) documented by Kahneman and Tversky (Kahneman and Tversky 1979). Much of the recent literature referred to as behavioral economics is based on this critique. The third tenet of Bayesianism, namely, that beliefs should be updated in accordance with Bayes's law, is almost unassailable from a normative viewpoint, when applied to the Grand State Space.⁸ It has, however, also been shown to be descriptively lacking. In fact, Tversky and Kahneman (1974) have shown that people often confound the conditional probability of A given B with that of B given A. The second tenet, namely, that one has probabilistic beliefs over anything uncertain, has been shown by Ellsberg (1961) to be an inaccurate description of people's behavior.⁹ Here we are concerned with the combination of the first two tenets, and argue that, taken together, they are far from compelling, not only as a description of human behavior or reasoning, but also as a normative ideal. More specifically, we accept the prevailing view in economic theory, according to which the most elegant way to think of uncertainty is to imagine the Grand State Space. But we do not see a rational way to assign probabilities to this space. Conversely, we have no quarrel with a Bayesian approach that assigns probabilities to a limited state space, especially if the latter has been encountered in similar problems in the past. It is only the combination of the first two tenets that we find problematic as explained below.

The main difficulty with assigning probability to the Grand State Space is that there is no information on which one can base the choice of prior beliefs. Any information that one may obtain, and that may help in the choice of a prior, should, according to the first tenet, be incorporated into the description of the Grand State Space. This means that the prior on this state has to specify beliefs one had before obtaining the information in question. Thus, the information may help one choose posterior beliefs (presumably according to Bayes's rule), but not prior beliefs.

Even when considering a state space that is not all-encompassing, there are many problems in which there is simply not enough information to sensibly generate probabilistic beliefs. In these problems one may expect people to exhibit behavior that cannot be summarized by a single probability measure. Moreover, when information is scarce, one may also reject the Bayesian approach on normative grounds, as did Knight (1921). This normative failure is related to the limitation discussed above: the Bayesian paradigm does not offer a theory of (prior) belief generation. It follows that,

⁷ See Mas-Collel et al. (1995, pp. 179–180).

⁸ When a more modest state space is considered, other approaches are possible. See Levi (1980).

⁹ For completeness, we describe Ellsberg's "paradox" in the Appendix.

even if one were convinced that one would *like* to be Bayesian, the Bayesian approach does not provide the self-help tools necessary to become Bayesian if one isn't already.

The normative failure of the second tenet of Bayesianism (Prior Probability) stimulated a search for alternative axiomatically-based models of decision theory. [Schmeidler \(1989\)](#) proposed a model with non-additive probabilities in which alternatives are ranked by Choquet expected utility. [Gilboa and Schmeidler \(1989\)](#) offered a model in which the decision maker entertains a *set* of priors, rather than a single one. The decision rule they axiomatized is the maxmin expected utility rule: every act is evaluated by the worst possible expected utility it could obtain, ranging over the various priors in the set, and this minimal expected utility is maximized by the decision maker.^{10, 11}

While the non-additive Choquet expected utility model and the maxmin expected utility model can be used to resolve Ellsberg's paradox (1961), they were not motivated by the need to describe observed behavior, but rather by the a-priori argument that the Bayesian approach is too restrictive to satisfactorily represent the information one has.

Consider the following example ([Schmeidler 1989](#)). You are faced with two coins, each of which is about to be tossed. The first coin is yours. You have tossed it, say, 1000 times, and it has come up Heads 500 times, and Tails 500 times. The second coin is presented to you by someone else, and you know nothing about it. Let us refer to the first coin as "known", and to the second as "unknown". Asked to assign probabilities to the known coin coming up Heads or Tails, it is only natural to estimate 50% for each, as these are the empirical frequencies gathered over a sizeable database. When confronted with the same question regarding the unknown coin, however, no information is available, and relative frequencies do not help estimate probabilities. But the second tenet of Bayesianism demands that both sides of the unknown coin be assigned probabilities, and that these probabilities add up to 1. Symmetry suggests that these probabilities be 50% for each side. Hence, you end up assigning the same probability estimates to the two sides of the unknown coin as you did for the two sides of the known coin. Yet, the two 50–50% distributions *feel* rather different. In the case of the known coin, the distribution is based on a good deal of information that supports a symmetric assessment while in the case of the unknown coin the same estimates are based on the absence of information. The Bayesian approach does not permit a distinction between symmetry based on information and symmetry based on lack of information. Importantly, one may specify the amount of evidence used to generate beliefs as part of the description of a "state", as dictated by the Grand

¹⁰ [Bewley \(2002\)](#) also axiomatized a set of priors, but in his model the decision maker has a partial preference order, whereby one act is preferred to another only if it has a higher expected utility, according to *each* possible prior.

¹¹ Both models are behavioral, in that they rely on in-principle observed preferences, and derive a representation thereof, which need not describe any actual mental process. Thus, [Schmeidler \(1989\)](#) suggests that people behave *as if* they maximize the Choquet integral with respect to a non-additive measure, and [Gilboa and Schmeidler \(1989\)](#) model decision makers that behave *as if* they are maximizing the minimal expected utility with respect to a set of priors. Both representations allow a naïve cognitive interpretation, but they are also compatible with other cognitive processes that may result in behavior that satisfies the relevant axioms.

State Space assumption, but this would generate the same problem when discussing the prior probability on the bigger state space.

One may embed this example in a decision problem, and predict choices as in Ellsberg's paradox. But it is important that the example above does not involve decision making. The point of departure of Schmeidler (1989) is not the descriptive failures of subjective EUT. Rather, it is what one might call a sense of cognitive unease with the manner that the Bayesian paradigm deals with absence of information. This cognitive unease points also to the normative failure of the Bayesian approach in this example. Even if one wished to become Bayesian, and even if one were willing to change one's choices so as to conform to the Bayesian paradigm, one must ignore the amount of information that was the basis of the prior beliefs.

The examples of Ellsberg and Schmeidler are simple in a way that might be misleading. These examples exhibit enough symmetries to suggest a natural prior via Laplace's "principle of insufficient reason". If one wished to become Bayesian, one could comfortably assign 50% probability to each color in Ellsberg's two-urn experiment, and, similarly, 50% to each side of the unknown coin in Schmeidler's example. In both cases, the 50–50% distribution is the only prior that respects the symmetry in the problem, and it is therefore a natural candidate for one's beliefs. Hence, considering these examples in isolation, one might conclude that, cognitive unease aside, it is fairly easy to become Bayesian even if one was not born Bayesian.

Unfortunately, most decision problems encountered in real life do not possess sufficient symmetries for the principle of insufficient reason to uniquely identify a prior. Consider, for example, the uncertainty about an impending war. One cannot seriously suggest that the relative frequency of wars in the past may serve as a good estimate of the probability of a war at the present. The occurrence of wars cannot be viewed as repeated identical and independent repetitions of the same experiment, and consequently the question of war is an example of uncertainty, rather than of risk. Applying Laplace's principle of insufficient reason would suggest that war has 50% probability, independent of the circumstances. But we know enough about war and peace to dismiss the symmetry between them. Indeed, we can reason at length about the likelihood of war, to have sufficient reason to reject the principle of insufficient reason.

There is an approach that holds that, by considering sufficiently detailed events, one can find "atomic" propositions on which the uniform prior is reasonable. We do not accept this point of view. There are well-known examples of the arbitrariness of Laplace's principle, dating back to Bertrand's paradox (suggested in 1888). More importantly, we hold that it is not rational to assign a prior probability of 50% to a proposition *only* because nothing is known about it. That is, if a Martian were asked about a probability of war, his knowing nothing about wars on Earth would not justify a 50–50% as rational belief for him. In fact, we believe that the only cases in which Laplace Principle of Indifference appears convincing involve at least one of two ingredients, absent in the war example: (i) the state space has certain symmetry properties which rely on a physical mechanism, such as the symmetry between two sides of a coin, or between different slots of a roulette wheel; (ii) the uniform prior can be justified as a limit belief when the state space is being divided according to an ever-finer partition. A theorem that justifies equal probabilities for "red" and "black"

in a roulette wheel spin on these grounds dates back to Poincaré (see a discussion and more citations in Machina, 2004). We do not know of any example in which Laplace's principle is justified that does not fall into one, or even both of these categories. And we hold that in the absence of any such structure, as in the case of the Martian assigning probabilities to war on Earth, it is simply irrational to choose a uniform prior rather than admitting ignorance.

To sum, a major failure of the Bayesian approach is that many real-life problems do not offer sufficient information to suggest a prior probability. In a small fraction of these problems there are symmetries that suggest a unique prior based on the principle of insufficient reason. But the vast majority of decision problems encountered by economic agents fall into a gray area, where there is too much information to arbitrarily adopt a symmetric prior, yet too little information to justifiably adopt a statistically-based prior.

Justification of beliefs by evidence offers a criterion for rationality that need not rank highly specified beliefs as more rational than less specified ones. While it is irrational to be agnostic and ignore evidence, it may also be irrational to hold probabilistic beliefs that are not grounded in evidence and are therefore arbitrary. As in the case of scientific work, which aspires to rationality, there are claims that science has established, and there are hypotheses about which science is silent. A high standard of rationality would require not only that one would not reject established facts, but also that one would not insist on claims that are unfounded.

4 Behavioral derivations of Bayesianism

The limitations of the Bayesian approach mentioned above are cognitive in spirit: they deal with the degree to which a mathematical model captures our intuition when reasoning about uncertainty. The standard approach in economics, however, would find these cognitive limitations hardly relevant. Modern economic theory follows logical positivism in viewing intentionally named concepts such as “utility” and “belief” as theoretical constructs that must be derived from observations. The revealed preference paradigm further holds that only observed choices are legitimate observations for economics, ruling out other sources of information such as introspection and subjective reports of preferences and likelihood judgments. Importantly, the derivations of subjective EUT by Ramsey, de Finetti, and Savage are consistent with this approach: they formulate axioms on (in-principle) observable choices, and show that these axioms are equivalent to a representation of choice by expected utility maximization according to a subjective probability measure.

Many economic theorists believe that the behavior of rational agents satisfies (say) Savage's axioms, and consequently that these agents behave as if they were maximizing expected utility with respect to an implicit Bayesian prior. Thus, a prevalent view is that rationality necessitates Bayesianism. Further, since the axioms identify the prior uniquely, this is often interpreted as revealed preferences being sufficient to elicit the beliefs of rational agents, namely, to measure their prior explicitly.

The claim that rational agents must follow Savage's axioms is often justified by arguments relying on "Dutch books" or on evolutionary reasoning. Both types of arguments suggest that individuals who violate the axioms will be "driven out" of the market.

In this section we argue that rational behavior need not imply Savage's axioms. We start by discussing two possible interpretations of a "preference order" and analyzing the axioms in light of these interpretations. We argue that under neither interpretation are the axioms a necessary implication of rationality. Further, we claim that eliciting a Bayesian prior from behavior is not always a viable option. Finally, we comment briefly on the weaknesses of "Dutch books" and evolutionary arguments.

4.1 Raw preferences and reasoned choice

We will address both the descriptive and the normative interpretation of Savage's axioms.¹² It will be more efficient to divide the discussion along different lines. Consider a binary relation representing preferences, or choices, as in Savage's theory. This relation can be interpreted in (at least) two ways. First, it might reflect *raw preferences*, namely, an instinctive tendency to prefer one alternative over another. For example, the decision maker may choose a piece of chocolate cake over a bowl of broccoli without thinking. Alternatively, the same binary relation might model *reasoned choice*, namely, choice that was arrived at by a process of reasoning. When asked to choose between a lottery that pays \$10 if the Democrats win both houses of Congress in the next election (and nothing otherwise) and a lottery that pays \$5 if the Republicans win both houses in the next election (and nothing otherwise), a decision maker is likely to think about the choices before deciding. Roughly, the decision maker exhibits raw preferences if she first acts, and then possibly observes her own act and stops to think about it. The decision maker is involved in reasoned choice if she first thinks, then decides how to act.

A descriptive interpretation of preferences in Savage's model may be either one of raw preferences or of reasoned choice. When describing reality, one must cope with the fact that in certain decision problems the decision maker acts before (or without) thinking, whereas in others she may reason her way to a decision. By contrast, a normative interpretation of Savage's theory deals with reasoned choice: if one attempts to convince a decision maker to change her decision(s), one would normally provide reasons to do so.

In the following two sub-sections we consider raw preferences and reasoned choice separately, and argue that Savage's axioms need not hold under either interpretation of "preferences".

4.2 Derivation of Bayesianism from raw preferences

There is ample evidence that raw preferences often do not satisfy Savage's axioms. The evidence ranges from Ellsberg's experiments, attacking a specific axiom, to the works

¹² The same discussion can be conducted in the context of any other behavioral axiomatization of Bayesianism. We choose to refer to Savage as his is, justifiably, the most well-known axiomatization.

of Kahneman and Tversky, questioning nearly every canon of rationality, explicit or implicit in the theory. Many economists take this evidence to mean that economic agents are not always rational, but they still hold that rational economic agents would satisfy the axioms. It is this claim that we take issue with. Hence, we do not resort to the “behavioral” argument that Savage’s axioms are violated in practice. Rather, we hold that raw preferences of an individual we would like to think of as “rational” may still deviate from Savage’s ideal. We propose two main arguments: first, we argue that in the face of uncertainty, the very notion of raw preferences may be in conflict with rationality. Second, we point out that the unobservability of the state space makes the meaning of Savage’s axioms vague, and renders impossible the elicitation of beliefs from choices.

4.2.1 Rational choices follow beliefs

We hold that economic agents, rational as they may be, only have raw preferences in certain domains. There are many questions of interest to economic analysis, where agents’ preferences simply do not exist prior to thinking. Moreover, there is nothing inherent in rationality that necessitates that raw preferences exist in all possible choice situations. To the contrary, it often appears more rational to base choices on explicit beliefs.

We do not claim that raw preferences are necessarily incompatible with rationality. People typically know which dish is tastier without having to deduce or compute their preference relation between dishes. In this sense, “preference needs no inference” (Zajonc 1980). Such preferences can certainly be compatible with utility maximization or any other behavioral standard of rationality. Moreover, in such cases one may indeed use an axiomatic derivation to justify and even calibrate a theoretical model. In the example of preferences between dishes, one may satisfy the axioms of completeness and transitivity without being aware of one’s own utility function. One may first observe one’s own choices between dish A and dish B, and then infer that her utility for the former is higher than for the latter. Inferring utility from observed choices is, indeed, the standard interpretation of a consumer’s utility function.

It is also possible that an agent would observe her own behavior, and subsequently learn what her implicit beliefs are. A driver who buckles up only when she drives on a highway might infer that she assigns a higher probability to the event of a serious accident when driving on the highway as compared to city driving, without explicitly estimating probabilities. However, it seems more rational to first consider the danger of accident, and then to decide whether to buckle up, rather than vice versa. Yet some routine behaviors may be rational without requiring explicit reasoning, mostly because they are familiar.

But many decision situations of interest are novel. For example, suppose Bob must make decisions whose outcomes depend on the possibility of a nuclear war in Asia in the next five years. Specifically, he is asked whether he prefers to get an outcome x if such a war erupts, and an outcome y otherwise, to, say, a sure outcome z . Bob cannot be expected to have a-priori preferences over these choices. In fact, it would appear highly irrational to have such preferences. Rather, he would stop and ponder, and only after assessing the probability of war can he meaningfully answer these preference

questions. Furthermore, this would seem the only rational way to generate preferences in this problem.

Those decision situations in which one has raw preferences from which beliefs may be inferred appear to be precisely those situations that are repeated often enough to allow the definition of beliefs by empirical frequencies. By contrast, in situations that are novel or unique, where one does not have sufficient information to generate a prior, one also does not have preferences that precede beliefs.

The revealed preference approach suggests that any observed choice manifests implicit preference. For example, if Bob has not decided to bet on the eruption of war, this will be taken to be a manifestation of certain beliefs about war, whether Bob thought about this decision or not. This default definition of “preferences” will guarantee that preferences exist, but it guarantees neither that they satisfy Savage’s axioms, nor that they appear “rational” in any intuitive sense.

4.2.2 *The unobservability of preferences*

The notion of “observability” in decision and economic theory allows some freedom of interpretation. Most theories in economics and in related disciplines have paradigmatic applications, which leave little room for multiple interpretations. For example, it is convenient to think of Savage’s axioms in the context of bets on a color of a ball drawn at random from an urn. In this context, the states of the world are clearly defined by the balls in the urn. Choices made contingent on the color of the ball drawn can be thought of as direct observations of a preference relation in a Savage-type model. However, most economic applications of subjective EUT do not have such a clearly defined state space and further, it is often not clear what state space the decision maker has in mind. The following example illustrates the difficulty.

Assume that we observe Mary deciding to quit her job and take another. This decision involves uncertainty and can be couched in a Savage-type model. A state of the world in such a model should specify the values of all variables that may be relevant to Mary’s choice. For instance, a state should be sufficiently detailed to determine Mary’s salary in a year’s time, both in her previous job and in the new job. One might also wish to include the expertise Mary might acquire on the job, her likely co-workers, the economic stability of her employer, her own job security, and a number of other variables, for any time horizon that one considers relevant to the problem. Moreover, the set of states of the world should also allow all possible causal relationships between Mary’s actions and these variables. For example, a reasonable model should be rich enough to express the possibility that with one employer Mary’s promotion is practically guaranteed, while with another it depends on her effort level. The nature of the causal relationship between effort and promotion is also subject to uncertainty, and should therefore also be specified by each state of the world.

These considerations give rise to two difficulties. First, the state space becomes large and complicated, and with it the set of conceivable acts defined on this state space in Savage’s model. Second, it is not at all clear which of the above uncertainties should be taken into account in our model. Hence, after observing Mary’s choice, we may construct many different Savage-type models in which her choice is modeled as a preference between two acts. But in each such model there will be many other pairs

of acts, the choice between which was not observed. Clearly, if the states of the world themselves are unobservable, one cannot hope to observe a complete binary relation between all the acts defined on these states.

A possible solution to the second problem would be to define an exhaustive state space, within which one may embed every conceivable state space that might be relevant to the problem. But such a solution aggravates the first problem. Moreover, it renders most pairwise choices inherently unobservable. To see this, imagine that one defines the set of outcomes to include all conceivable consequences, over any time horizon. One then proceeds to define states as all possible functions from acts to outcomes. This would result in an exhaustive, canonical state space. Next, one must define all the conceivable acts (Savage 1954): all possible functions from states to outcomes. Over this set of conceivable outcomes one assumes that a complete binary relation is observable, and that the observed choices would satisfy Savage's axioms. But such a relation cannot be observable even in principle. In this states are functions from actual acts to outcomes, and conceivable acts are functions from these states to the same set of outcomes. Thus the set of conceivable acts is by two orders of magnitude larger than the set of acts that are actually available in the problem. This implies that the vast majority of the binary choices assumed in Savage's model are not observable, even in principle.

Savage's theory has a clear observable meaning in experiments involving simple set-ups such as balls drawn out of urns, but in many economic applications of EUT the state space is not directly observable, and hence Savage's behavioral axioms do not have a clear meaning in observable terms. There are two implications of this. First, in many situations of interest it will be impossible to elicit prior beliefs from raw preferences. Second, the claim that rational agents have preferences that satisfy Savage's axioms becomes ill-defined. Since there are many possible Savage-style models in which the decision problem can be couched, the claim should be re-stated as holding that *there exists* a Savage-style model in which a rational agent satisfies the axioms. But this existential claim is quite weak, and perhaps vacuous.

4.3 Derivation of Bayesianism from reasoned choice¹³

We now turn to the interpretation of a preference relation as describing reasoned choice. We argue that also under this interpretation rationality need not imply compliance with Savage's axioms, and that elicitation of beliefs from choices is not always possible.

4.3.1 Reasoned choice need not be complete

The completeness axiom is typically justified by necessity: one must make a decision, and whatever one chooses will be viewed as the preferred act. This argument seems to apply to observed preferences. Indeed, if one defines preference by observations,

¹³ See Shafer (1986) for related arguments.

the completeness axiom is rather weak, but when we consider reasoned choice, there is no compelling argument for completeness.

To see this point more clearly, consider first the case of transitivity. If there is a reason to prefer f to g , and if there is a reason to prefer g to h , then these two reasons may be combined to provide a reason to prefer f to h . The transitivity axiom may actually be viewed as a reasoning axiom, providing an argument, or a justification for a certain preference. It can thus be used to infer certain preferences from others. Similarly, three other axioms in Savage's framework can be viewed as templates for reasoning that a certain act should be preferred to another. Specifically, Savage's sure thing principle suggests that, if preference hold between two acts that are equal on a given event, then the same preference should hold between another pair of acts, which are also equal on that event and are unchanged off that event. Savage's monotonicity axiom states that, if we replace the outcome that an act f yields on an event A by a better outcome, the modified act should be preferred to the original one. The fourth axiom of Savage that can be viewed as a reasoning axiom relates preferences for betting on an event A versus betting on an event B , where the two pairs of acts involved vary the notion of "good" and "bad" outcomes involved in the bet. All these axioms can help reach a decision about what one's preferences between two acts should be, given one's preferences between other pairs of acts.¹⁴

The same cannot be said of the completeness axiom. The completeness axioms states that (reasoned) choice should be defined between any two acts f and g , but it provides no help in finding reasons to prefer f to g or vice versa.

If one views Savage's axioms as conditions on raw preferences, the completeness axiom may be mentioned as a half-axiom barely worth mentioning. Completeness in this set-up is one of two requirements in Savage's first axiom, which is well-accepted in consumer theory and in choice under certainty. But if the Savage axioms are viewed as conditions for reasoned choice, the completeness axiom plays an altogether different role: it is contrasted with all the rest. The completeness axiom defines the question, namely, what are the reasoned preferences between pairs of acts, and all the rest are part of the answer, that is, potential reasons that may come to bear in determining preferences between particular pairs of acts.

How should we model rational choice when there are no obvious reasons to assume preferences are complete? One possibility is to adhere to purely reasoned choice, namely to choice that can be justified. This will result in relaxation of the completeness axiom, as was suggested by [Bewley \(2002\)](#). Another approach is to model explicitly the choice that will eventually be made, and thus to retain the completeness axiom, but to relax some of the other axioms that might not be as compelling under uncertainty. The resulting model attempts to describe the choice that will eventually be made, even if it cannot be fully justified. (See [Schmeidler 1989](#), and [Gilboa and Schmeidler 1989](#).)

¹⁴ Savage employed three additional axioms. The first only requires that preferences be non-trivial, and is needed to fix a unique probability measure of the decision maker, while the other two axioms are continuity axioms. While such axioms can also be viewed as part of reasoned choice, namely, as helping one complete one's own preferences, for the purposes of the present discussion we focus on those that are more fundamental from a conceptual viewpoint.

The two approaches agree that reasoned choice need not be complete. Both admit that rationality does not imply Savage's axioms.

4.3.2 Computational complexity

When decision problems do not present themselves to the decision maker with a clearly defined state space, the generation of all relevant states involves insurmountable computational difficulties. For example, assume that a contract involves n binary conditions. The state space used to analyze it would need to have 2^n states. Imagining these states and reasoning about their likelihood may prove a daunting cognitive task. A similar argument can be made for finding regularities in given databases. In fact, some of these problems can be proved to be “difficult” in a precise sense (see [Aragonès et al. \(2005\)](#)). Because certain problems are computationally so complex as to prohibit a solution by the most advanced computers in existence, it also appears unlikely that rational economic agents can solve them. Just as rational economic agents do not play chess optimally, they cannot be assumed to solve problems that computer scientists cannot solve.

4.3.3 Other difficulties

The elicitation of beliefs from reasoned choice encounters two additional problems that were discussed also for raw preferences. First, the observability problem means, for reasoned choice, that it is not clear which state space the agent has in mind. Resorting to a canonical state space may not reflect the agent's reasoning, and it would also mean that defining one's preference relation requires a large degree of hypothetical reasoning. Second, reasoned choice might contradict other axioms of Savage, beyond the completeness axiom. While these axioms are themselves reasons for a particular preference, they may generate conflicting preferences. The only known algorithm that would guarantee a resolution of these conflicts in a way that satisfies the axioms would be to select a probability and a utility function and to make decisions in accordance with EUT. This, however, brings us back to the task of specifying a prior directly, and would be a failure of the attempt to derive prior beliefs from choices.

4.4 Dutch books and evolution

The “Dutch book” or “money pump” arguments suggest that a violation of a certain axiom would lead one to lose money with certainty. For example, having cyclical (strict) preference, say, of a over b , b over c , and c over a , one is exposed to a sequence of trades, where one pays a small amount for exchanging one's choice for a preferred alternative, ending up with the original choice but with less money. The possibility of such a “money pump” would presumably lure a shrewd bookie who would drive the intransitive agent into bankruptcy and out of the market.

Dutch book arguments of this type rely on the assumption that the decision maker in question is wrong about the model she uses. In the example above, a preference of a over b should mean that the outcome a , which presumably specifies all that matters

to the decision maker's well-being, is preferred in her eyes to another outcome b , which is also understood to be a final outcome. If, by contrast, these outcomes are only transient states, at which additional decisions need to be made, with potentially additional trades, the decision maker was not employing a correct model in the first place: she ignored the dynamic aspect of the trade involved. Such a decision maker does not satisfy a very basic condition of rationality, namely, that she is aware of the "right" model, that is, the one used by the external observer.

An analysis of Dutch book arguments was recently proposed by Hájek (2008). We mention in passing that such an irrationality of the decision maker is also likely to be corrected as a result of experience in real life: confronted with a Dutch book and trading a for b , b for c , and c for a , for a small sum each time, people would notice that something is wrong way before they are driven into destitution. The irrationality that Dutch books expose has to do with a fundamental error of modeling on the decision maker's part. Such an error is far from our standard of rationality irrespective of one's preferences, and cannot prove the irrationality of the latter.

A more serious argument is offered by evolutionary reasoning. Here one does not posit a bookie who exploits a particular agent's preferences, but relies on a long-term process, in which various decision situations naturally suggest themselves to the agent, and where, asymptotically, violation of certain axioms might lead to sub-optimal results. The latter are assumed to drive the preferences in question into extinction, relying on processes of replication, say, under imitation.

Evolutionary arguments offer a powerful theory about the type of rationality that one might expect to observe. When rationality is independently defined, we may pose the question, whether a certain evolutionary process does or does not lead to rational behavior. But evolutionary arguments cannot serve as a definition of rationality per se. Moreover, evolutionary arguments are compelling in the presence of sufficient repetitions of choice situations in similar environments. But these are precisely the situations in which one may have sufficient information to generate probabilistic beliefs. For example, it stands to reason that evolution has prepared the human mind to make rational decisions about the likelihood of meat being rotten; it is less obvious that evolution has prepared us to make rational decisions in face of global warming. More generally, it is precisely when one confronts novel situations, and does not have sufficient information for the generation of such beliefs, that the evolutionary arguments appear the weakest.

5 Conclusion

Economic research tends to assume that agents who face uncertainty are Bayesian, that is, that they behave as if they intended to maximize expected utility with respect to a subjective probability measure, which is updated according to Bayes's rule when new information is obtained. This assumption has been challenged experimentally by Ellsberg (1961). Indeed, in recent years alternative models, such as Choquet expected utility (Schmeidler 1989) and maxmin expected utility with multiple priors (Gilboa and Schmeidler 1989) have been used in economic applications ranging

from finance and game theory to labor economics, international trade, and macroeconomics.¹⁵

Many economists are willing to accept the need for non-Bayesian models in order to improve the accuracy of economic theories. They view non-Bayesian decision making as a form of “bounded rationality” that should be sprinkled into a theory to make it descriptively more accurate, as one may provide better predictions by introducing hyperbolic discounting, mental accounting, or other biases and mistakes that are being incorporated into “behavioral” economic models. Our claim in this article is different. We do not focus on descriptive failures of the Bayesian approach, but on its *normative* inadequacies. We challenge the Bayesian approach as the gold standard of rationality.

While we do not offer a definition of “rationality”, the preceding discussion suggests a distinction among three levels of rationality, as applied to choice or to belief. At the highest level choices and beliefs are *reasoned*, namely: the individual can defend and justify them by appeal to objective evidence. At the next level decisions and beliefs are not justified by evidence, but are not in conflict with evidence. This category may include raw preferences that are robust to the reasoning process. Finally, irrational beliefs are those that are contradicted by evidence. Similarly, irrational choices are those that are in conflict with reason, and that may well be changed as a result of the reasoning process.

The behavioral failure of Savage’s theory in the case of Ellsberg’s experiments is not “irrational” in this sense. There is substantial evidence that individuals prefer to bet with known probabilities rather than with unknown ones, even after the sure thing principle has been explained to them. We have argued, however, that taking the higher standard of rationality, that is, of reasoned beliefs and reasoned choice, one may end up with non-Bayesian beliefs. It is sometimes *more* rational to admit that one does not have sufficient information for probabilistic beliefs than to pretend that one does.

Much of this article was devoted to the claim that rationality of beliefs does not necessitate a probabilistic representation thereof. This argument was made on cognitive as well as on behavioral grounds, where behavior could mean raw preferences or reasoned choice, and could be interpreted descriptively or normatively. As argued in the introduction, the converse is also false: probabilistic beliefs may not be rational in any intuitive sense of the world. These two claims converge on the need to model the way beliefs are formed. A theory of belief formation will help us determine which Bayesian beliefs are rational, and will also indicate when it is rational to hold Bayesian beliefs in the first place.

¹⁵ Dow and Werlang (1992) were the first to apply such models to finance, explaining why an individual may wish neither to buy nor to sell an asset for a range of prices. Epstein and Wang (1994, 1995) and Epstein and Schneider (2003) have extended the theory to a dynamic setting and used it to explain stock markets booms and crashes. Mukerji (1998) applied the theory to explain incompleteness of contracts. Epstein and Miao (2002) and Alonso (2005) use these models to explain the “home bias” phenomenon. Ozaki and Nishimura (2002) use the theory to explain why uncertainty, as opposed to risk, decreases an unemployed individual’s search time. Hansen and Sargent (2001, 2003, 2006) used these models, and ideas from robust control theory, to develop a new approach to macroeconomic theory.

Appendix: Ellsberg's paradox

One of Ellsberg's "paradoxes" is the following. There are two urns. Urn 1 contains 100 balls, each of which is known to be either red or black, but you have no information about how many of the balls are red and how many are black. Urn 2 contains 50 red balls and 50 black balls. A **red bet** is a bet that the ball drawn at random is red and a **black bet** is the bet that it is black. In either case, winning the bet, namely, guessing the color of the ball correctly, yields \$100. First, you are asked, for each of the urns, if you prefer a red bet or a black bet. For each urn separately, most people say that they are indifferent between the red and the black bet.

Then you are asked whether you prefer a red bet on urn 1 or a red bet on urn 2. Many people say that they would strictly prefer to bet on urn 2, the urn with known composition. The same pattern of preferences is exhibited for black bets (as, indeed, would follow from transitivity of preferences given that one is indifferent between the betting on the two colors in each urn). That is, people seem to prefer betting on an outcome with a known probability of 50% than on an outcome whose probability can be anywhere between 0 and 100%.

It is easy to see that the pattern of choices described above cannot be explained by expected utility maximization for any specification of subjective probabilities. Such probabilities would have to reflect the belief that it is more likely that a red ball will be drawn from urn 2 than from urn 1, *and* that it is more likely that a black ball will be drawn from urn 2 than from urn 1. This is impossible because in each urn the probabilities of the two colors have to add up to one. Thus, Ellsberg's findings suggest that many people are not subjective expected utility maximizers. Moreover, the assumption that comes under attack is not the expected utility hypothesis per se: *any* rule that employs probabilities in a reasonable way would also be at odds with Ellsberg's results. The questionable assumption here is the second tenet of Bayesianism, namely, that all uncertainty can be quantified in a probabilistic way. Exhibiting preferences for known versus unknown probabilities is incompatible with the second tenet of Bayesianism.

Ellsberg's paradox and Schmeidler's two-coin example are simple illustrations of the distinction Knight made between "risk" and "uncertainty". These examples demonstrate that Savage's axioms may fail in descriptive theories. Specifically, Savage imposes the "sure thing principle" (axiom P2), which states the following. Should two possible acts, f and g , yield precisely the same outcomes if an event A does not occur, then the preference between f and g should only depend on their outcomes given A . That is, if we were to modify these acts outside of A , in such a way that they are still equal to each other, then the preferences between the modified acts should be the same as between the original ones. This axiom appears very compelling: if the two acts are anyway equal off A , why should one care what they are equal to? It seems natural to ignore the values of the acts when they are equal, and to focus on the event in which differences among them might emerge. Yet, it can be seen that the behavior exhibited in Ellsberg's paradox violates this axiom as would a decision maker who prefers to bet on a known coin rather than on an unknown one in Schmeidler's example.

To see why P2 is violated in these examples, let the state space have four states, each specifying the outcome of a draw from the known urn I as well as the outcome of

a draw from the unknown urn II (even though only a draw from one urn is possible). For example, the first state might be denoted I-R;II-R, suggesting that a draw from urn I would result in a red ball, and so will a draw from urn II. Similarly, we have three more states: I-R;II-B I-B;II-R and I-B;II-B. Respecting this order of the states, each act is associated with a vector consisting of four outcomes. Let f be the act “betting on Red in urn I”, which is associated with the vector of outcomes (1,1,0,0) (that is, winning in the first state, I-R;II-R, and the second state, I-R;II-B, and losing in the other two states). Let g be the act “betting on Red in urn II”, whose vector of outcomes is (1,0,1,0). Acts f and g are equal on the event { I-R;II-R I-B;II-B } (to 1 in the first state and to 0 in the second). P2 would suggest that, if we were to modify both f and g on this event to 0 and 1 (instead of 1 and 0), preferences between them would not change. But in this case the modified f is associated with (0,1,0,1), that is, it is the act “betting on Black in urn II”, whereas the modified g , (0,0,1,1), is the act “betting on Black in urn I”, and the preferences between these are reversed. Intuitively, the modification of act f changed it from a risky act, whose distribution is known, to an uncertain (or “ambiguous”) act, whose distribution is not known. In the case of act g , the same modification changed an uncertain act into a risky one. This asymmetry may explain why people violate P2 in this example, and often feel comfortable with this violation even when the logic of the axiom is explained to them.

Acknowledgments Our thinking on these issues was greatly influenced by discussions with many people. In particular, we wish to thank Edi Karni, Marion Ledwig, Dan Levin, Stephen Morris, Peter Wakker, the editors, and two anonymous referees for comments and references.

References

- Alonso, I. (2005). *Ambiguity in a two-country world*. Princeton: Mimeo, Princeton University.
- Aragonès, E., Gilboa, I., Postlewaite, A. W., & Schmeidler, D. (2005). Fact-free learning. *American Economic Review*, 95, 1355–1368.
- Bewley, T. (2002). Knightian decision theory: Part I. *Decisions in Economics and Finance*, 25, 79–110. (Working paper, 1986).
- Carnap, R. (1952). *The continuum of inductive methods*. Chicago: University of Chicago Press.
- Cho, I.-K., & Kreps, D. M. (1987). Signaling games and stable equilibria. *The Quarterly Journal of Economics*, 102, 179–222.
- de Finetti, B. (1937). La Prevision: Ses Lois Logiques, Ses Sources Subjectives. *Annales de l'Institut Henri Poincaré*, 7, 1–68.
- Dow, J., & Werlang, S. R. C. (1992). Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica*, 60, 197–204.
- Epstein, L. G., & Miao, J. (2003). A two-person dynamic equilibrium under ambiguity. *Journal of Economic Dynamics and Control*, 27, 1253–1288.
- Epstein, L. G., & Schneider, M. (2003). Recursive multiple priors. *Journal of Economic Theory*, 113, 32–50.
- Epstein, L. G., & Wang, T. (1994). Intertemporal asset pricing under knightian uncertainty. *Econometrica*, 62, 283–322.
- Epstein, L. G., & Wang, T. (1995). Uncertainty, risk-neutral measures and security price booms and crashes. *Journal of Economic Theory*, 67, 40–82.
- Ellsberg, D. (1961). Risk, ambiguity and the savage axioms. *Quarterly Journal of Economics*, 75, 643–669.
- Ghirardato, P. (2002). Revisiting savage in a conditional world. *Economic Theory*, 20, 83–92.
- Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics*, 18, 141–153. (Working paper, 1986).
- Hájek, A. (2008). Dutch book arguments. In Anand Paul, P. Prasanta, & P. Clemens (Eds.), *The Oxford handbook of rational and social choice*. Oxford: Oxford University Press.

- Hansen, L. P., & Sargent, T. J. (2001). Acknowledging misspecification in macroeconomic theory. *Review of Economic Dynamics*, 4, 519–535.
- Hansen, L. P., & Sargent, T. J. (2003). Robust control of forward-looking models. *Journal of Monetary Economics*, 50, 581–604.
- Hansen, L. P., & Sargent, T. J. (2006). *Robustness*. Princeton: Princeton University Press (forthcoming).
- Jeffrey, R. (2004). *Subjective probability: The real thing*. Cambridge: Cambridge University Press.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–291.
- Knight, F. H. (1921). *Risk, uncertainty, and profit*. Boston: Houghton Mifflin.
- Levi, I. (1980). *The enterprise of knowledge: An essay on knowledge, credal probability, and chance*. Cambridge: MIT Press.
- Lindley, D. (1965). *Introduction to probability and statistics from a Bayesian viewpoint*. Cambridge: Cambridge University Press.
- Machina, M. J., & Schmeidler, D. (1992). A more robust definition of subjective probability. *Econometrica*, 60, 745–780.
- Machina, M.J., & Schmeidler, D. (1995). Bayes without bernoulli: Simple conditions for probabilistically sophisticated choice. *JET*, 67, 106–128.
- Mailath, G., Masahiro, O.-F., & Andrew, P. (1993). On belief based refinements in signaling games. *Journal of Economic Theory*, 60, 241–276.
- Mas-Colell, A., Michael, W., & Jerry, G. (1995). *Microeconomic theory*. New York: Oxford Press.
- Mukerji, S. (1998). Ambiguity aversion and the incompleteness of contractual form. *American Economic Review*, 88, 1207–1232.
- Ramsey, F. P. (1931). Truth and probability. In R. B. Braithwaite (Ed.), *The foundation of mathematics and other logical essays*. New York: Harcourt, Brace and Co.
- Rostek, M. (2010) Quantile maximization in decision theory. *The Review of Economic Studies*, 77, 339–371.
- Savage, L. J. (1954). *The foundations of statistics*. New York: John Wiley and Sons. (Second addition in 1972, Dover).
- Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57, 571–587. (Working paper, 1982).
- Shafer, G. (1986). Savage revisited. *Statistical Science*, 1, 463–486.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124–1131.
- Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist*, 35, 151–175.