AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Is It Always Rational to Satisfy Savage's Axioms?

Itzhak Gilboa, Andrew Postlewaite, and David Schmeidler

3-09

June, 2009

The Foerder Institute for Economic Research
and
The Sackler Institute of Economic Studies

Is It Always Rational to Satisfy Savage's Axioms?¹

by

Itzhak Gilboa², Andrew Postlewaite³, and David Schmeidler⁴

April 2009

This note argues that, under some circumstances, it is more rational not to behave in accordance with a Bayesian prior than to do so. The starting point is that in the absence of information, choosing a prior is arbitrary. If the prior is to have meaningful implications, it is more rational to admit that one does not have sufficient information to generate a prior than to pretend that one does. This suggests a view of rationality that requires a compromise between internal coherence and justification, similarly to compromises that appear in moral dilemmas. Finally, it is argued that Savage's axioms are more compelling when applied to a naturally given state space than to an analytically constructed one; in the latter case, it may be more rational to violate the axioms than to be Bayesian.

Can probabilities reflect ignorance?

Will the US President six years hence be a Democrat? The Bayesian approach requires that we be able to quantify this uncertainty by a single number; we should be able to state that our subjective belief for this event is, say, 62.4% or 53.7%. Many people feel that they do not have sufficient information to come up with such an accurate probability estimate. Moreover, some people feel that it is more rational not to assign a probabilistic estimate for such an event than to assign one. Choosing one probability number in the interval $[0,1]$ would be akin to pretending that we know something that we don't.

Consider another example. There is a semi-popular talk at your university, titled, "Cydophines and Abordites". You are curious and may listen to the talk, after which you'll probably know what these terms mean, what is known about them and so forth. However, before the talk you have no idea what the terms mean. In fact, you do not even know what discipline they belong to. For all you know, these could be designating enzymes, grammatical structures in an ancient language, or Abelian groups. You are asked whether all cydophines are abordites. Obviously, you have no

¹ We wish to thank Eddie Dekel, Dov Samet, and Peter Wakker for discussions and comments.

² Tel-Aviv University, and HEC, Paris. tzachigilboa@gmail.com

³ University of Pennsylvania. apostlew@econ.upenn.edu

⁴ The Ohio State University and Tel-Aviv University. schmeid@tau.ac.il

idea. But if you are Bayesian, you should have probabilistic beliefs about this fact. How would you be able to come up with the probability that all cydophines are abordites?

You may be tempted to assign a 50%-50% prior to the claim that all cydophines are abordites. But classical discussions of Laplace's Principle of Indifference (or Principle of Insufficient Reason) show that the seemingly neutral 50%-50% prior doesn't lead very far. For instance, if there are such concepts as pre-cydophines or semi-cydophines, are they, too, abordites with probability of 50%? And what about super-abordites? Should we perhaps assign equal probabilities to the four possibilities: cydophines are a sub-class of abordites; abordites are a sub-class of cydophines; cydophines and abordites are disjoint; cydophines and abordites are logically independent? And after we are done with this question, what should be the probability that cydophines are red? Should it be uniform over the color words in the language, or over the visible spectrum? Or should we first condition this proposition on cydophines having a color as an attribute to begin with?

We claim that it is irrational to assign a probability to cydophines being abordites. There is no logical inconsistency in choosing any number to be the probability of the proposition in question, but it appears irrational to choose arbitrarily such a number and insist that it is *the* probability of the proposition. The Bayesian approach is lacking because it is not rich enough to describe one's degree of confidence in one's assessments.⁵ For any probability question it requires a single probability number as an answer, excluding the possibility of replies such as "I don't know" or "I'm not so sure". A paradigm of rational belief should allow a distinction between assessments that are well-founded and those that are arbitrary.

What is a rational decision?

The Bayesian approach could therefore be viewed as an elegant but imperfect method for representation of uncertainty, one among many to be used depending on the application. Indeed, this is the way that it is viewed by many in diverse fields such as statistics, philosophy, and computer science.⁶ However, within economic theory the Bayesian approach is the sole claimant to the throne of rationality.⁷ The most important reason is probably the axiomatic foundations of subjective expected utility maximization. Building on works of Ramsey (1931) and de Finetti (1937), as well as by von Neumann and Morgenstern (1944), Savage (1954) provided the most compelling axiomatic derivation of this theory. Starting with very abstract objects, and posing a few compelling axioms, he showed that a complete preference relation that satisfies the axioms can be represented as maximizing expected utility relative to a subjective probability measure. Thus, under seemingly weak assumptions one may conclude that we should formulate our beliefs in terms of a Bayesian prior and make decisions so as to maximize the expectation of a utility function relative to this prior.

⁵ See Knight (1921), Ellsberg (1961), as well as Shafer (1986) and Schmeidler (1989).

⁶ See Carnap (1952), Lindley (1965), Levi (1980), and Jeffrey (2004).

⁷ See, for example, the standard text by MasColell, Whinston, and Green (1995), and, for a recent contribution, Al-Najjar and Weinstein (2008).

From a normative viewpoint, the theory appears to be very persuasive. Moreover, many believe that there is no mathematical result in the entire corpus of the social sciences that compares to Savage's theorem in terms of elegance and generality, as well as conceptual and mathematical depth.

Savage did not intend his theory to apply to every conceivable source of uncertainty. But if the theory is so enticing, why shouldn't we always adopt it, at least as a normative goal? Indeed, much of economic theory does precisely this. If asked, what is the "rational" way to make decisions in face of the uncertainty about cydophines and abordites, many economic theorists would suggest the Bayesian approach as the only rational decision making procedure. Our difficulty is that, despite the beauty of Savage's result, we still find it problematic to assign a probability, say, 72.3%, to the cydophines being abordites and argue that this is rational.

To explain our notion of rational choice, consider the following scenario. You are a public health official who must make a decision about immunization of newborn babies. Specifically, you have a choice of including another vaccine in the standard immunization package. This vaccine will prevent deaths from virus A. But it can cause deaths with some probability. The exact probabilities of death with and without the vaccine are not known. Given the large numbers of babies involved, you are quite confident that some fatalities are to be expected whatever your decision is. You will have to face bereaved parents and perhaps lawsuits. Will it be rational for you to pick prior probabilities arbitrarily and make decisions based on them?

We argue that the answer is negative. What would then be the rational thing to do, in the absence of additional information? Our main point is that there may not be any decision that is perfectly rational. There is a tension between the inability to justify any decision based on statistical data, scientific research, and logical reasoning on the one hand, and the need to make a decision on the other. This tension is well recognized and it is typically resolved in one of two ways. The first is the reliance on default choices. If the choices that can be rationally justified result in an incomplete preference relation, a default is used to make decisions where justified choice remains silent. For example, the medical profession suggests a host of "common practices" that are considered justified in the absence of good reasons to deviate from them. The second approach is to avoid defaults and to use a complete preference relation that incorporates caution into the decision rule. For example, dealing with worst-case scenarios, which is equivalent to a maxmin approach, can be suggested as a rational decision rule in the face of extreme uncertainty.

To consider another example, suppose that the decision maker is the US administration who has to decide whether to wage war against a country that is suspected of producing nuclear weapons. There is uncertainty about the state of the technology of that country as well as about its intentions. Military and political science experts are consulted, but their views differ. There seems to be no agreed upon, or objectively justifiable answer to the question, "what is the probability that the country in question will possess operational nuclear weapons within one year?" This question is important. Moreover, different probability values will lead to

different decisions. Will it be rational for the administration to assume a value, say, 90%, and make the best decision based on this value?

Again, we argue that it would not be rational to do so. One can hardly defend such a weighty choice on the basis of an arbitrary probability, chosen so as to satisfy Savage's axioms. If the best decision for a probability of 90% is different from the decision for, say, 20%, many would feel that rationality precludes the possibility of choosing the former value and behaving as if it were known. As in the medical example, there are two standard ways out. One is to assume that the status quo of peace should be adopted unless one has good reasons to discard it. Thus, no war is the default decision. Alternatively, one can adopt a worst-case analysis, admitting that one does not have a precise probability estimate, but arguing for caution in the face of uncertainty.

Both approaches may be viewed as less than perfectly rational. Indeed, using a default choice legitimates phenomena that are generally regarded as boundedly rational, such as a status-quo bias. On the other hand, the maxmin rule violates Savage's axiom P2, which appears eminently rational. Our point is that satisfying Savage's axioms is also not perfectly rational, for the reasons mentioned above. That is, in some situations *there may not be a perfectly rational choice at all*.

We reject the view that rationality is a clear-cut, binary notion that can be defined by a simple set of rules or axioms. There are various ingredients to rational choice. Some are of internal coherence, as captured by Savage's axioms. Others have to do with external coherence with data and scientific reasoning. The question we should ask is not whether a particular decision is rational or not, but rather, whether a particular decision is more rational than another. And we should be prepared to have conflicts between the different demands of rationality. When such conflicts arise, compromises are called for. Sometimes we may relax our demands of internal consistency; at other times we may lower our standards of justifications for choices. But the quest for a single set of rules that will universally define *the* rational choice is misguided.

Analogy between rationality and morality

We find that the question, "what is the rational thing to do?" bears structural similarity to the question, "what is the moral thing to do?" We explain this analogy below. Obviously, some readers may accept our view when applied to rationality and reject it when morality is concerned; or vice versa; or accept our views on both, but find the analogy weak. With these caveats we offer the analogy between morality and rationality in the hope that it will clarify our view of the latter.

When dealing with rationality as well as with morality, one may adopt an a priori, axiomatic approach, subscribing to a set of rules that by definition dictate the "right" thing to do, moral or rational. This would be the case if one adopts, say, the Ten Commandments as the definition of moral conduct, or Savage's axioms as the definition of rational behavior. But in both cases one may also adopt a different view,

according to which the general principles are only approximations, which can be further refined in light of particular examples.

We think of the definition of morality and of rationality as an act of modeling. The modeler attempts to capture people's preferences for an ideal mode of behavior. These preferences often do not coincide with people's actual behavior, that is, with revealed preference. For example, people may feel that it is immoral to lie, but sometimes may find themselves lying. Similarly, people may wish to be dynamically consistent, yet sometimes find themselves yielding to temptation. When dealing with the definition of morality or of rationality, we take a normative point of view, rather than a descriptive one: we attempt to model the behavior that people *would like* to exhibit, rather than the behavior they actually do exhibit. Still, these are data that need to be captured by models: people's preferences, intuitions, and desires are given, and they should be described by the model.

People's preferences about their behavior, that is, the type of behavior they would like to exhibit, exist at two levels (at least): rules that apply to single actions and consistency principles that deal with the comparison of actions. In the context of morality, a consistency principle might be "equal treatment of equals". This principle does not say anything about the preferred behavior in any particular case. It is consistent with malevolent intentions, as long as these are fairly directed at everyone. But when it is coupled with one's benevolence towards some (perhaps one's self), it yields a code of behavior that we may find acceptable, or at least a step in the right direction. Importantly, people have preferences both regarding rules, such as doing good rather than doing evil, and about the consistency principles, such as "equal treatment of equals". The preferences for rules may sometimes be in conflict. So may be the preferences for consistency principles. Worse still, rules and consistency principles may interact to generate contradictions that make some moral dilemmas non-trivial. Indeed, it is often not clear what "the moral thing to do" is, and one often has to make compromises, discuss more or less moral choices, and so forth.

Savage's axioms are consistency principles, analogous to "equal treatment of equals". In isolation, these principles do not put any constraints on one's beliefs. Hence, they are insufficient for a definition of rationality. A definition of rationality that does not impose additional constraints on beliefs beyond Savage's consistency principles would be analogous to a definition of morality that satisfies itself with equality principles, but remains silent on which deeds are moral in and of themselves. Thus we are led to ask, what conditions we expect rational behavior to satisfy. The question can be addressed both to tastes and to beliefs: one may ask whether it is rational to make decisions according to a particular utility function, or according to a given probability measure. Our focus here is on the second question.

Which additional constraints should we impose on rational belief? For example, "rationality" should mean, to most people, some coherence with scientific data. It is irrational to believe that smoking is not detrimental to one's health. No violation of Savage's axioms is involved in behaving according to such a belief: if one starts with zero prior probability that smoking is dangerous, one ends up with zero posterior

probability for this event. But it is irrational to hold such beliefs in the face of evidence.

It appears that a minimal requirement of rationality is that one not hold beliefs that are contrary to objectively available data, coupled with logical, statistical, or mathematical reasoning. A higher standard of rationality demands that one only subscribe to beliefs that can be so justified. According to this notion, it is irrational to behave as if one had good reasons to hold certain beliefs where one actually does not.⁸

When we identify several consistency principles, as well as several rules for justification of belief, we should be prepared to encounter contradictions. Again, one may draw an analogy to questions of morality. We may have a strong sense that it is moral to give some money to a panhandler on the street. We may also feel that all equally poor people should be entitled to the same level of support. Coupled together, we may find that it is impossible to follow both the rule and the consistency principle. In such situations most of us seem to be psychologically prepared to live with compromises, sometimes following the rule but violating the consistency principle, sometimes the other way around.

We maintain that the question of rationality is similar. Given various rules of rationality on the one hand, and a collection of consistency principles on the other, one may not be surprised to find occasional contradictions. In this case the question is not “what is *the* rational thing to do?” but “what is more rational to do in this instance?” Correspondingly, the answer we give may be subjective and imperfect. The question of rationality becomes murkier than we would have liked it to be. Indeed, even the concept of “justification” of beliefs is quantitative and fuzzy. Yet, it seems to us more rational to admit that such trade-offs exist rather than to stick to consistency principles alone, totally ignoring the demand that beliefs be justified.

Compromises

The previous sections provided examples in which we find it more rational to admit ignorance than to pretend that probabilities can be assigned to all propositions. In this case, how do we respond to a Savage questionnaire? Which of Savage’s axioms will we be willing to sacrifice, and when?

There are two axioms that are natural candidates to be violated: the completeness axiom and P2 (often referred to as the “Sure Thing Principle”). For instance, in the cydophines-abordites example we may simply refuse to express preferences over Savage acts defined over a state space involving these unknown terms. We may restrict our preferences to those that we can justify in some reasonable sense, and we remain silent about many others.⁹

The allegedly behavioral elicitation of beliefs from choices a la Savage assumes a notion of “a situation repeated under the same conditions”: one needs to assume that

⁸ This point has been argued and elaborated in Gilboa, Postlewaite, and Schmeidler (2004).

⁹ See Bewley (2002).

all pairs of acts are compared over the same state space, that is, that the state does not change from one choice to the other. In many situations, this is highly hypothetical. Moreover, as pointed out elsewhere,¹⁰ if we wish to define the state space in a way that allows for all possible causal relationships, we end up with a Savage model in which the vast majority of act pairs cannot be compared: the set of “conceivable acts” is larger by two orders of magnitude than the set of acts actually available to the decision maker. In short, violating the completeness axiom need not be as theoretically costly as often assumed in economics.

Consider again the question of the party of president of the US six years hence. As opposed to cydophines and abondites, one may argue that this problem is relevant to many economic questions, and violating completeness in this context does restrict the power of our models: decisions will eventually be made, and a model with incompleteness leaves part of the story untold. In this case one may, with a heavy sigh, accept completeness and decide to relax P2.¹¹ It is not an easy decision to make, if one aspires to be rational. But the alternative is not so enticing either. Given the choice between the rule that says, “I will base my probabilities on evidence and calculation” and the consistency principle embodied in P2, there does not seem to be a unique “rational” decision.

It is important, however, that accepting the consistency principles in their full strength and ignoring any other rules is not sufficient for rationality. Internal coherence of beliefs is important, but so is external coherence: having coherent beliefs that have nothing to do with evidence and data cannot be considered rational. One may cling to all of Savage’s axioms, but then one would have to admit that Savage’s result is an impossibility result: it shows that the seemingly compelling consistency principles lead to a result that is patently counter-intuitive.

Where do the states originate?

Whether we accept a certain rule or a consistency principle may depend on the application. We need not commit to a principle across all decision problems or reject it in all: we may find it applicable in some but less in others. Pushing the analogy between rationality and morality a little further, rules and principles of morality can also be qualified by the type of application.

In this context let us consider the potential violation of completeness and of the Sure Thing Principle in Ellsberg’s two experiments (Ellsberg, 1961). The experiments are well-known, and so is their analysis. We briefly repeat it here to highlight the difference between them.

In the single-urn experiment, there are 90 balls, of which 30 are red, and the rest are blue or yellow.

¹⁰ An argument that we have spelled out in Gilboa and Schmeidler (1995, 2001), and Gilboa, Postlewaite, and Schmeidler (2004).

¹¹ See Schmeidler (1989), Gilboa and Schmeidler (1989), among many others.

	Red	Blue	Yellow
Red	1	0	0
Blue	0	1	0
Not-Red	0	1	1
Not-Blue	1	0	1

Suppose that a decision maker is uncertainty averse, and therefore prefers “Red” to “Blue” and “Not-Red” to “Not-Blue”. This is a clear violation of P2: since “Red” and “Blue” are equal on Yellow, changing both of them from 0 to 1 on Yellow should result in the same preference, but this would mean that “Not-Blue” should be preferred to “Not-Red”, and not the other way around.

In the two-urn experiment there is a known urn (urn I), with 50 black and 50 red balls, and an unknown urn (urn II), with 100 balls that are black or red. The decision maker is allowed to choose *both* the urn and the color, and bet on a random draw from the chosen urn having the specified color. Suppose that the decision maker is again uncertainty averse, and therefore prefers each of the bets on the known urn to each of the bets on the unknown urn. This clearly implies that the decision maker is not probabilistically sophisticated, that is, has preferences that cannot be described by preferences over distributions (relative to a single probability measure): no subjective probability measure can assign both Red and Black in the unknown urn probabilities that are strictly less than 50%.

But it is not so immediate to verify that these preferences violate P2. To see this, one first has to construct the state space, such that each state specifies the outcome of a draw from the known urn as well as a draw from the unknown urn.¹² The matrix one gets would be as follows:

	I-B; II-B	I-B; II-R	I-R; II-B	I-R; II-R
I_B	1	1	0	0
I_R	0	0	1	1
II_B	1	0	1	0
II_R	0	1	0	1

¹² One may include in the description of the state also the number of black (and red) balls in urn II, but there is no substantial loss in suppressing this source of uncertainty.

In this matrix, an act “I_B” means, “bet on black out of urn I”, “II_R” – “bet on red out of urn II”, etc. Each state is a function from the set of urns {I,II} to the set of colors {B,R}, and the four states are all such functions.

Given this decision matrix, the violation of P2 becomes evident: consider the two acts I_B and II_R. They coincide on the event $A = \{ \text{“I-B; II-R”}, \text{“I-R; II-B”} \}$. According to P2, if we change the values of both I_B and II_R on this event from (1,0) to (0,1), the preferences between them should not change. But this results in I_B becoming II_B, and II_R becoming I_R, and uncertainty averse preferences rank I_R above II_R, not below it.

But there is something artificial in the state space analysis in this example. The states are not naturally given in the problem; they were analytically constructed to fit into the mold of a decision matrix in which a state “resolves all uncertainty”. In fact, the states in this problem will never be revealed to the decision maker: she has to choose to bet on urn I *or* on urn II. Thus, at the end of the game, she will either have the partition

$$\{ \{ \text{“I-B; II-R”}, \text{“I-B; II-B”} \}, \{ \text{“I-R; II-R”}, \text{“I-R; II-B”} \} \}$$

or the partition

$$\{ \{ \text{“I-B; II-B”}, \text{“I-R; II-B”} \}, \{ \text{“I-B; II-R”}, \text{“I-R; II-R”} \} \}$$

but not both. The state of the world that truly obtains will never be revealed to the decision maker.

Furthermore, the event A used to apply Savage’s P2 in this example is a highly contrived event: it reads “Either a ball from the known urn is black and a ball from the unknown is red, or vice versa”. This event will never be observed by the decision maker: she will never know whether it has or has not occurred. By contrast, in the single-urn experiment the event playing this role was Yellow. It was naturally given in the description of the problem, and, importantly, an event whose truth value could be verified post-hoc.

The set of acts over which Savage assumes complete preferences also differs qualitatively in the two experiments. In the single urn experiment, assuming only three states¹³, we have to consider 8 acts, all of which are easily imagined. It makes sense to offer the decision maker the 8 bets and asked that they be ranked. By contrast, in the two urn experiments there are only four acts that are actually available, but 16 acts that are “conceivable”. That is: considering the minimal state-space model in which the problem can be couched, one has to consider states as functions from acts to outcomes, and then to consider all the acts that are functions from this (analytically constructed) state space to outcomes. This set is by two orders of magnitude larger than the set of acts that are actually available in the original

¹³ Obviously, this is incompatible with Savage’s axiom P6. Our arguments are only strengthened when one takes into account the complications induced by infinite state spaces, measurability constraints, etc.

problem. Moreover, adding more acts, say, over yet another urn, would require a re-definition of the state space, and with it – of the set of conceivable acts.

Consider again the question of the party of the President of the US six years hence, and compare it to the result of a toss of a fair coin, say, at election time, so that the two sources of uncertainty do not differ in terms of the time of their resolution. Assume that the decision maker tells us that she prefers to bet on either side of the coin to either Democrat or Republican. We may tell her that her preferences cannot be described by a probability measure. Suppose that she is indifferent to this derogatory remark, and we try to convince her that she also violates P2. We will have to tell her to imagine a state space in which each state specifies what will be the President's party in six years, *and* what will be the outcome of the coin. This is a little strange, because the two bets will not occur simultaneously. Then we will ask the decision maker to observe that on the event “the president is a Democrat and the coin comes up Head or the President is a Republican and the coin comes up Tail”, two choices gave (1,0), but were they to give her (0,1), she would find that...

Suppose that the decision maker is sufficiently sophisticated and mathematically oriented to see the point. And suppose that she wishes to satisfy P2. But if the cost is that she has to commit to the probability of the next President being a Democrat being a particular number, she may decide that certain violations of P2 may be less irrational than making up priors and taking them seriously.

The problem is accentuated in situations where the decision maker suspects that her choices are causally related to external circumstances. For instance, assume that the decision maker is the President of the US, who has to decide whether to wage war on another country, or whether to save a major bank facing the risk of bankruptcy. The lives or livelihood of many people are at stake. Each possible choice has an uncertain outcome. But one would never know what would have been the outcome of the choices that were not made. Similarly, if the decision maker is an economic agent who has to decide whether to start a new business in industry A or industry B, she will know how successful is the business she ended up engaging in, but she will not know how well off she would be had she made another decision. The Bayesian approach calls for the generation of prior beliefs about outcomes given each possible choice. To justify this demand by Savage's axioms, one has to consider an analytically constructed state space as above. But the decision maker will choose only one of her options. Therefore, as in the two-urn experiment described above, she will never be able to tell which state obtains. Importantly, this feature is inherent to the problem: each state has to describe the outcome of all acts, while only one act will actually be chosen.

This discussion highlights an implicit axiom in Savage's model: the claim that all uncertainty can be represented by the states of the world. This assumption per se is costless: one can always imagine a canonical state space, where each state specifies all the truth values of all propositions that one may even be interested in. But when this assumption is coupled with the other Savage's axioms it becomes far from innocuous. With a very large state space we also have a larger (indeed, very much

larger) set of acts. Having a complete preference that satisfies all of Savage's axioms may be reasonable in a "small", naturally-given state space. It is a much more demanding proposition when the state space is analytically constructed so as to describe all that might ever matter to the decision maker.

References

- Al-Najjar, Nabil I. and Jonathan Weinstein (2008), "The Subjective Approach to Ambiguity: A Critical Assessment", this volume.
- Bewley, Truman (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.
- de Finetti, Bruno (1937), "La Prevision: Ses Lois Logiques, Ses Sources Subjectives", Annales de l'Institute Henri Poincare, 7: 1-68.
- Ellsberg, Daniel (1961), "Risk, Ambiguity and the Savage Axioms", Quarterly Journal of Economics, 75: 643-669.
- Gilboa, Itzhak, Andrew Postlewaite, and David Schmeidler, "Rationality of Belief", 2004.
- 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, "Case-Based Decision Theory", The Quarterly Journal of Economics, 110 (1995) 605-639.
- Gilboa, Itzhak, and David Schmeidler, A Theory of Case-Based Decisions. Cambridge University Press, 2001.
- Jeffrey, Richard (2004), Subjective Probability: The Real Thing, Cambridge: Cambridge University Press.
- Knight, Frank H. (1921), Risk, Uncertainty, and Profit. Boston, New York: Houghton Mifflin.
- Levi, I. (1980), The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance. Cambridge: MIT Press.
- Lindley, Dennis (1965), Introduction to Probability and Statistics from a Bayesian Viewpoint. Cambridge: Cambridge University Press.
- Mas-Colell, Andreu, Michael Whinston, and Jerry Green (1995), Microeconomic Theory. New York: Oxford Press.
- 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 addition in 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.
- von Neumann, John, and Oskar Morgenstern (1944), Games and Economic Behavior. New York: Wiley.