

Bayesian Probability

Patrick Maher
University of Illinois at Urbana-Champaign

November 24, 2007

ABSTRACT. “Bayesian probability” here means the concept of probability used in Bayesian decision theory. It is usually identified with the agent’s degrees of belief but that interpretation is unsatisfactory because it severely restricts the contexts in which Bayesian decision theory is correctly applicable. A satisfactory interpretation is obtained by taking Bayesian probability to be an explicatum for inductive probability given the agent’s evidence.

1 Introduction

Bayesian decision theory, in its usual form, asserts that rational choices maximize expected utility. Since expected utility is defined in terms of probability and utility, it follows that an understanding of Bayesian decision theory requires an interpretation of these concepts of probability and utility. In this paper I focus on probability; I will call the concept of probability used in Bayesian decision theory “Bayesian probability.”

We can stipulate what Bayesian probability is, but not every stipulation gives a satisfactory interpretation of Bayesian decision theory. Some stipulations are unsatisfactory because they make Bayesian decision theory inapplicable or incorrect in many situations. The question with which this paper is concerned is: What interpretation of Bayesian probability is satisfactory, in the sense that it makes Bayesian decision theory both applicable and correct?

The usual view is that Bayesian probabilities are the agent’s degrees of belief; I argue that this is unsatisfactory because it severely restricts the contexts in which Bayesian decision theory is correctly applicable. But before getting to that, I discuss three other interpretations of Bayesian probability, one of which is the interpretation I favor.

2 Physical probability

Suppose I am about to toss a coin and I tell you that this coin either has heads on both sides or else has tails on both sides (though I don’t say which); if I ask you to state the probability that it will land heads, there are two natural answers: (i) $1/2$; (ii) either 0 or 1 but I don’t know which. Both answers are right in some sense, though they are incompatible, so “probability” in ordinary

language must have two different senses. I call the sense of “probability” in which (i) is right *inductive probability*; I call the other sense, in which (ii) is right, *physical probability*.

Inductive probability is relative to evidence and independent of facts about the world; to know the inductive probability of the coin landing heads you need to know what the evidence is, not what is on the coin. Physical probability is the opposite, it isn’t relative to evidence and does depend on facts about the world; to know the physical probability of the coin landing heads you need to know what is on the coin, not what evidence is available.

Physical and inductive probability are both ordinary language concepts, that is, senses of the word “probability” in ordinary language. There are also philosophical theories of probability that attempt to explicate one or other of these ordinary language concepts. In particular, frequency and propensity theories of probability are best regarded as attempts to explicate the concept of physical probability. However, we can discuss physical probability without committing ourselves to any explication of it, and that is what I will do here.

I will now consider:

Proposal 1 *Bayesian probability is physical probability.*

Physical probabilities are relative to something; on the view I have defended (Maher 2008), they are relative to a type of experiment. Thus a full articulation of Proposal 1 would need to specify what experiment type is to be used. (In discussions of the frequency theory, this is called “the reference class problem.”) I won’t pursue that issue because, however it is resolved, the following problems will remain.

First, in many decisions problems the states¹ don’t have physical probabilities. For example, suppose John wants to propose to Mary if she loves him and not if she doesn’t; here the states must specify whether Mary loves John but there isn’t a physical probability that Mary loves John.

Second, even when the states do have physical probabilities, agents often don’t know what they are and so can’t use them to guide their decisions. For example, suppose I must decide whether to bet that a coin will land heads, and I know that the coin has either two heads or two tails, though I don’t know which; here I can take the states to be “the coin lands heads” and “the coin lands tails,” and these have physical probabilities, but I don’t know the values of these physical probabilities.

Thus adoption of Proposal 1 would have the result that Bayesian decision theory can only be applied to a very narrow class of problems, namely, those in which the states have physical probabilities and the agent knows the values of these physical probabilities. For this reason, Proposal 1 is unsatisfactory.

¹ In decision theory, the “states” are ways the world might be that, together with the act chosen, determine the consequence obtained by the agent (Maher 1993, 1–5).

3 Inductive probability

Since physical probability is unsatisfactory for our purposes, it is natural to next try the other probability concept of ordinary language, inductive probability. Before doing that, I will say more about this concept.

First, inductive probability isn't the same thing as degree of belief. To see this, suppose I claim that a theory H is probable in view of the available evidence. The reference to evidence shows that this is a statement of inductive probability. If inductive probability were degree of belief then I could justify my claim by proving that I have a high degree of belief in H . However, this doesn't really justify my claim; to justify it I need to cite features of the available evidence that support H . It follows that inductive probability isn't degree of belief. Of course, any sincere intentional assertion *expresses* the speaker's degrees of belief; the point I am making is that assertions of inductive probability aren't assertions *about* the speaker's degrees of belief.

Let an *elementary probability sentence* be a sentence which asserts that a specific hypothesis has a specific numeric value; for example, "The probability of H given E is 0.5," where H and E are specific sentences. I say that a probability concept is *logical in Carnap's sense* if all true elementary sentences for it are analytic.² Every probability function whose numeric values are specified by definition is logical in Carnap's sense, so there demonstrably are probability concepts that are logical in Carnap's sense. In particular, the function c^* in (Carnap 1950) is logical in this sense, since its numeric values are fixed by its definition. Similarly, every c_λ in Carnap's (1952) continuum of inductive methods is logical in Carnap's sense. On the other hand, physical probability isn't logical in Carnap's sense because there are elementary sentences of physical probability whose truth value depends on contingent facts, such as what is on a coin.

The truth value of elementary sentences of inductive probability doesn't depend on facts about the world (as physical probability does) or on the speaker's psychological state (as degree of belief does). There also don't appear to be any other contingent facts that it depends on. We thus have good reason to conclude that inductive probability is also logical in Carnap's sense. Since the only other ordinary language concept of probability is physical probability, and it isn't logical in Carnap's sense, we can define inductive probability as the concept of probability in ordinary language that is logical in Carnap's sense.

Concepts of ordinary language are often vague and inductive probability is no exception. For example, the inductive probability that there is intelligent life in the Andromeda galaxy, given what I know, has no precise numeric value. Nevertheless, some inductive probabilities have precise numeric values; in my example of the double-sided coin, it is plausible that the inductive probability of the coin landing heads has the precise value $1/2$.

² This terminology is motivated by the characterization of logical probability in Carnap (1950, 30). For a refutation of Quine's criticisms of analyticity, see Carnap (1963, 915–922).

For further discussion of the concept of inductive probability, see (Maher 2006a). I will now consider:

Proposal 2 *Bayesian probability is inductive probability given the agent's evidence.*

This has the advantage over Proposal 1 that it allows Bayesian decision theory to be applied in some contexts in which physical probabilities either don't exist or are unknown; that is because inductive probabilities can exist and be known when physical probabilities either don't exist or are unknown. My example of the double-sided coin is a case of this kind.

However, Proposal 2 does have a serious drawback. As noted above, inductive probabilities often lack precise numeric values. Bayesian decision theory in its standard form requires numeric probabilities for the states and so Proposal 2 has the result that decision theory in its standard form is only applicable in the rather special situations in which the inductive probabilities of the states, given the agent's evidence, have precise numeric values.

To deal with this drawback, one might think of representing inductive probabilities by sets of probability functions and then modifying Bayesian decision theory to say that a rational choice must maximize expected utility relative to at least one probability function in the set. But this won't work because inductive probabilities not only often lack precise numeric values, they also lack precise upper and lower boundaries to their vagueness, whereas sets of probability functions determine precise upper and lower boundaries for each probability.³ In the next section I will present a better way of dealing with the vagueness of inductive probability.

4 Explication of inductive probability

Carnap described a philosophical methodology that he called *explication*; it consists in identifying a precise concept that can be used in place of a given vague concept, at least for certain purposes. The vague concept is called the *explicandum* and the corresponding precise concept is called the *explicatum*. Carnap (1950, 7) stated the requirements for an explicatum as follows:

1. The explicatum is to be *similar to the explicandum* in such a way that, in most cases in which the explicandum has so far been used, the explicatum can be used; however, close similarity is not required, and considerable differences are permitted.
2. The characterization of the explicatum, that is, the rules of its use (for instance, in the form of a definition), is to be given in an *exact* form, so as to introduce the explicatum into a well-connected system of scientific concepts.

³ Numeric probability values are real numbers in the interval $[0,1]$, hence a set of them is bounded both above and below, from which it follows that the set has a least upper bound and a greatest lower bound, which I refer to informally as its "precise boundaries."

3. The explicatum is to be a *fruitful* concept, that is, useful for the formulation of many universal statements (empirical laws in the case of a nonlogical concept, logical theorems in the case of a logical concept).
4. The explicatum should be as *simple* as possible; this means as simple as the more important requirements (1), (2), and (3) permit.

Suppose we take as our explicandum the vague concept of inductive probability. I would choose as explicatum a conditional probability function p whose values are specified by definition. It follows that p is logical in Carnap's sense and hence is, in this respect, like inductive probability. In order for p to be a good explicatum in other respects, we need to choose the definition of p suitably. So, for example, if the inductive probability of H given E has a precise numerical value, we would normally want $p(H|E)$ to be equal to that value. If the inductive probability of H_1 given E is greater than that of H_2 given E (which can be true even if these probabilities lack precise numeric values), then we would normally want to have $p(H_1|E) > p(H_2|E)$. One can easily specify other similar desiderata.

I can now state the proposal I favor; it is:

Proposal 3 *Bayesian probability is an explicatum for inductive probability given the agent's evidence.*

In order to apply this proposal to a decision problem we must have an explicatum for inductive probability given the agent's evidence. However, that explicatum can be very limited in scope; what is required is only that for each state S we have a number which we accept as an explicatum for the inductive probability of S given the agent's evidence. These numbers may be obtained in any of the ways in which people now come up with "reasonable" numbers for the probabilities of the states, ranging from unaided intuition to formal statistical models. The difference here consists more in how we interpret these numbers than in how they are obtained.

Acceptance of Proposal 3 implies a reinterpretation of Bayesian decision theory. We will no longer think of it as making a claim about the vague pre-theoretic concept of rational choice but rather will construe it as proposing an explicatum for that concept. This explicatum is maximization of expected utility, where expected utility is calculated using a probability function that is an explicatum for inductive probability and a utility function that is an explicatum for the value of the possible consequences to the agent.

Proposal 3 shares the advantages of Proposal 2, since it doesn't use physical probabilities. It also avoids the defect of Proposal 2, since the explicatum for inductive probability always has, as a matter of stipulation, numerically precise values. We don't even need to complicate Bayesian decision theory, as is done by those who futilely attempt to accommodate vagueness by using sets of probability functions; on this approach we are guaranteed a unique probability function and so can use Bayesian decision theory in its standard form.

Granted, the recommendations of Bayesian decision theory on this conception have a certain artificial precision about them, but that is the nature of explication and it helps us better understand vague ordinary language concepts, as I have argued elsewhere (Maher 2007).

I have already mentioned that the methodology of explication and my conception of logical probability both come from Carnap (1950). In addition, the concept I call “inductive probability” appears to be the same as what Carnap (1950) called “probability₁” and took as his explicandum. So Proposal 3 is merely an application of concepts from Carnap (1950). Nevertheless, this proposal differs greatly from the logical interpretation of probability commonly attributed to Carnap. According to this common view, Carnap in (1950) believed there is a unique probability function c such that $c(H|E)$ is the degree of belief in H that is rational for a person whose total evidence is E . Proposal 3 differs from this in the following ways:

1. It distinguishes between inductive probability and an explicatum for it; no such distinction is made in the conception commonly attributed to Carnap.
2. It doesn’t postulate a unique numerically precise probability function. Inductive probability is unique but not numerically precise; an explicatum for it is numerically precise but not unique.
3. It doesn’t say anything about rationality. When I said that inductive probability and explicata for it are logical in Carnap’s sense, I was making a semantic claim, not an epistemic or normative one, as you can verify by checking my definition of “logical in Carnap’s sense.” There is no interesting sense in which rationality requires a person’s degrees of belief to agree with either inductive probability or an explicatum for it (Maher 2006a, 189–191).

Thus, although Proposal 3 is based on concepts articulated long ago by Carnap, it is a new proposal that has not even been considered, much less refuted, by contemporary authors.

5 Degree of belief

For the past 40 years or so, almost all decision theorists have accepted:

Proposal 4 *Bayesian probabilities are the agent’s degrees of belief.*

I used to accept this proposal too (Maher 1993); at that time I hadn’t conceived of Proposal 3 and didn’t appreciate the force of the criticisms of Proposal 4 that I am now going to make.

5.1 Three criticisms

First criticism: In Bayesian decision theory, as usually formulated, the probabilities of the states are precise numbers. On the other hand, people's degrees of belief often lack numerically precise values. Thus Proposal 4 has the result that Bayesian decision theory, as usually formulated, often cannot be applied.

Advocates of Proposal 4 usually respond to this criticism by saying that the agent's degrees of belief can be represented by the *set* of probability functions that are consistent with the degrees of belief that the person has. They then need to modify Bayesian decision theory so that it only assumes a set of probability functions, not a unique function. This modification of Bayesian decision theory has been developed differently by different authors, but all versions agree that rationality requires choosing an act that maximizes expected utility relative to at least one probability function in the set of probability functions that represents the person's degrees of belief (Levi 1986; Kaplan 1996).

However, real people's degrees of belief not only lack precise numeric values, they also lack precise numeric upper and lower boundaries, whereas a set of probability values necessarily has precise boundaries. Thus, the move to sets of probability functions, and the consequent complication of Bayesian decision theory, doesn't avoid the problem, which is this: If we adopt Proposal 4 then Bayesian decision theory is often not applicable because the agent's degrees of belief are insufficiently precise.

Second criticism: Real people often have degrees of belief that are inconsistent with the probability calculus (Kahneman et al. 1982). So if we represent the agent's degrees of belief by the set of probability functions that are consistent with the person's degrees of belief, that set will often be empty; in this case, we get a definite set but, since it is empty, it doesn't permit the application of Bayesian decision theory.

Third criticism: Even if the agent's degrees of belief satisfy the laws of probability, they may be unrelated to the agent's evidence; in such cases, using Proposal 4 has the result that Bayesian decision theory gives wrong judgments about what choices are rational. For example, in 2003 there was no good evidence that Iraq had weapons of mass destruction but George W. Bush was nevertheless practically certain that it did. If we suppose that Bush's degrees of belief were representable by a non-empty set of probability functions then (with some other assumptions) Bayesian decision theory using Proposal 4 will say that Bush was rational to invade Iraq, but this is surely not correct.

The first two criticisms show that if Proposal 4 is adopted then Bayesian decision theory isn't applicable to many realistic decision problems. The third criticism shows that, even when the theory can be applied, it may give incorrect results. Consequently, Proposal 4 is unsatisfactory.

Proposal 3 isn't subject to any of these criticisms. When we adopt Proposal 3, it doesn't matter if the agent's degrees of belief are vague because we use an explicatum for inductive probability that is precise. It doesn't matter

if the agent’s degrees of belief are inconsistent with the laws of probability because we use an explicatum for inductive probability that satisfies those laws. And it doesn’t matter if the agent’s degrees of belief are unrelated to the agent’s evidence because we use an explicatum for inductive probability that takes appropriate account of the evidence.

5.2 The normative defense

My criticisms of Proposal 4 rest on facts that are mostly familiar to the many decision theorists who endorse that proposal. The usual response is to say that Bayesian decision theory is a normative theory, not a descriptive one. For example, they say it isn’t a relevant objection that real people’s degrees of belief are often inconsistent with the probability calculus; Bayesian decision theory is about what is rational and rational degrees of belief are consistent with the probability calculus.

This defense sounds plausible if you don’t think about it too carefully but in fact what is asserted is neither true nor relevant, as I will now show.

5.2.1 Not true

We saw that if Proposal 4 is adopted then Bayesian decision theory is often inapplicable to real agents because their degrees of belief are insufficiently precise. The normative defense is that *rational* people’s degrees of belief are sufficiently precise and thus Bayesian decision theory applies to rational people. However, there is no sense of the term “rational” in which rationality requires a person’s degrees of belief to either have precise numeric values or precise upper and lower boundaries. For example, if we take rational degrees of belief to be ones that it would be advisable to adopt then, since there are usually better things to do with one’s time than to try to make one’s degrees of belief precise, rationality in this sense doesn’t require such precision. If rational degrees of belief are ones that are appropriately related to the evidence then, since the appropriate relation is often vague with no precise boundaries, rationality in this sense also doesn’t require degrees of belief to have precise numeric values or boundaries (Maher 2006b, 144–147).

We also saw that if Proposal 4 is adopted then Bayesian decision theory doesn’t apply to real agents whose degrees of belief are inconsistent with the laws of probability. For brevity, I will call degrees of belief that are consistent with the laws of probability *coherent*; my criticism then is that Proposal 4 makes Bayesian decision theory not apply to agents whose degrees of belief are incoherent. This is the issue to which adherents of Proposal 4 have devoted most of their attention, producing a vast literature which attempts to show, in one way or another, that *rational* degrees of belief must be coherent. On some ways of understanding it, this claim is also false.

1. If “rational degrees of belief” are ones that are advisable to adopt, then they need not be coherent. That is because it would take a lot of time

and effort to make one's degrees of belief coherent and there are usually better things to do with one's time.

2. If “rational degrees of belief” are ones that are appropriately related to the evidence, and if a degree of belief that is a precisification of the corresponding inductive probability counts as appropriately related to the evidence, then again it is false that rational degrees of belief must be coherent (Maher 2006b, 143–144).
3. If “rational degrees of belief” are ones that *equal* the corresponding inductive probability, then it is true that rational degrees of belief are coherent, since inductive probabilities are. However, this conception of “rational degree of belief” means that two people with the same evidence must have the same degrees of belief in order to be rational, which is denied by most advocates of Proposal 4.

5.2.2 Not relevant

We need to distinguish between whether a normative theory correctly *describes* real agents and whether it *applies* to real agents. Laws against murder don't correctly describe murderers but they do apply to them; otherwise they would be pointless. My criticism of Proposal 4 was that in many common situations it prevents Bayesian decision theory from correctly *applying* to real agents. The normative response is that we can't expect a normative theory to correctly *describe* real agents; that is true but not a relevant response to the objection.

Advocates of Proposal 4 sometimes suggest that Bayesian decision theory applies to real agents by telling them they ought to revise their degrees of belief so as to conform to the assumptions of the theory. However, if that is what the theory says, and if “ought to” means this is advisable, then the theory is false, since we usually have better things to do with our time than to make our degrees of belief precise and consistent with the laws of probability. If “ought to” just means that this would be nice in some sense, without implying anything about what the agent should do, then the theory is giving no advice to the agent, which is precisely what I am objecting to. Furthermore, on this conception the advice Bayesian decision theory gives to real agents is very unspecific (it doesn't say in what way this conformity should be achieved) and in particular it doesn't answer the question one expects decision theory to answer, namely, which of the options in a decision problem are rational choices. Some subjectivists seem to believe that Bayesian decision theory cannot answer this question, but with Proposal 3 it can.

5.3 Do what you like?

There is a very simple decision theory which I call the *do-what-you-like theory*; it says that rational choices are ones the agent prefers. This theory seems

obviously unsatisfactory, and has never been advocated by anyone so far as I know, but what is wrong with it?

In my view, there are two things wrong with it. First, in many decision problems the agent is undecided and has no preference, in which case the theory cannot be applied. Second, agents often have irrational preferences, in which case the theory gives incorrect results.

An advocate of the do-what-you-like theory might respond that the theory is about what is rational, so the fact that it isn't correctly applicable to real agents is irrelevant; rational agents always prefer the rational choices, so the theory is correct for rational agents. This is the normative defense of the do-what-you-like theory. Since the do-what-you-like theory is unsatisfactory and this defense would vindicate it, there must be something wrong with this defense. In my view, what is wrong with it is:

1. In some senses of the term "rational," it is false that rational agents always prefer rational choices. In particular, if a rational agent is one who always does what is advisable to do, then it isn't rational to prefer the rational choice in every decision problem, since the cost of acquiring all those preferences would outweigh the benefit; most of these decision problems will never arise.
2. The defense isn't a relevant answer to the objection. A useful decision theory needs to be correctly applicable to agents who don't prefer the rational choice in a decision problem, not only to those who do. The normative defense only observes that decision theory need not correctly *describe* real agents, which is not in dispute.

A stubborn defender of the do-what-you-like theory might respond that the theory does give advice to irrational agents, namely, it tells them that they ought to become rational agents; in particular, they ought to acquire a preference for the options that it is rational to choose. Since the do-what-you-like theory is unsatisfactory and this defense would vindicate it, there must be something wrong with this defense. In my view, the things wrong with it are:

1. If what is being claimed is that agents would be well advised to do whatever it takes to acquire a preference for the rational choice in any decision problem, then the claim is false, as I have already noted. If that isn't what is being claimed, then the theory isn't in fact giving any advice as to what irrational agents should do.
2. Second, if advice is being given, it is not only wrong but also very unspecific, since it doesn't say which preferences are rational; thus it doesn't tell us what we expect a decision theory to tell us, namely, which are the rational choices in a decision problem.

My criticisms of the do-what-you-like theory are the same as my criticisms of Proposal 4. Therefore, if you accept my criticisms of the do-what-you-like theory, you should also accept them as sound criticisms of Proposal 4.

Conversely, if you don't accept my criticisms of Proposal 4, then you may find the do-what-you-like theory even more to your liking than Bayesian decision theory with Proposal 4. After all, both theories can be defended in the same way and the do-what-you-like theory is simpler.

5.4 Representation theorems

I will refer to Bayesian decision theory with Proposal 4 as *subjective Bayesian decision theory*. The most sophisticated defense of this theory has been to prove that if a person's preferences satisfy certain axioms, such as transitivity, then there exists a probability function p and a utility function u such that the person's preferences maximize expected utility relative to p and u . A result of this sort is called a *representation theorem*. If we assume that a rational person's preferences must satisfy the axioms, and that the function p is a measure of the person's degrees of belief, then a representation theorem shows that the preferences of rational agents conform to subjective Bayesian decision theory. Ramsey (1926) was the first to prove a theorem of this kind; others include Savage (1954), Maher (1993), and Joyce (1999).

Representation theorems don't address any of the three criticisms that I have raised against Proposal 4, they just shift the issue from degrees of belief to preferences. If the criticisms are restated in terms of preferences they apply just the same, as I will now show. Since the issues here are essentially the same as before, I will be brief.

First criticism: Representation theorems assume that the agent has definite preferences regarding all the acts in a very large class but real agents don't have such complete preferences. Therefore, real agents needn't satisfy the assumptions of a representation theorem. Furthermore, rationality doesn't require a person to have such complete preferences, so even rational agents needn't satisfy the assumptions of a representation theorem.

In (1993, 19–21) I acknowledged the correctness of this but thought I could get around it by arguing that a rational person's preferences can be represented by the *set* of complete preference orderings that are consistent with the person's incomplete preferences and satisfy the axioms of a representation theorem. However, since even a rational person's preferences may be vague, there need not be any such set. In addition, this response is irrelevant since the objection is that the theory doesn't *apply* to real agents, not that it fails to *describe* real agents.

Second criticism: Real people's preferences often violate other axioms besides the one that requires them to be complete and so they aren't consistent with any complete preference ordering that satisfies the axioms. In this case, there is a set of complete preference orderings that satisfies the axioms and is consistent with the agent's preferences but it is empty and hence useless for representing the agent's preferences. Once again, the theory doesn't apply.

Third criticism: Even if the agent's preferences do satisfy all the axioms of a representation theorem, they needn't be rational. In that case, there will

be a p and u such that the agent's preferences all maximize expected utility relative to p and u , but some of these preferences will be irrational. Subjective Bayesian decision theory makes incorrect recommendations for such agents.

In addition to not resolving any of the three criticisms, representation theorems have additional liabilities of their own. They are complex, require dubious axioms that aren't implied by Bayesian decision theory, and are motivated by a dubious reduction of degree of belief to preference. By using Proposal 3 instead, none of this is necessary.

6 Conclusion

The history of science shows that methodological views which now appear to have little merit are sometimes nearly universally endorsed in a discipline for several decades. Examples include the mechanical philosophy in seventeenth century natural science and behaviorism in early twentieth century psychology. I believe the dominance of Proposal 4 in late twentieth century Bayesianism will soon be seen as a similar episode. In this paper I have tried to hasten the end of this episode by showing that Proposal 4 is inadequate and articulating a viable alternative to it.

References

- Carnap, Rudolf. 1950. *Logical Foundations of Probability*. University of Chicago Press. 2nd ed. 1962.
- . 1952. *The Continuum of Inductive Methods*. University of Chicago Press.
- . 1963. Replies and systematic expositions. In *The Philosophy of Rudolf Carnap*, ed. Paul Arthur Schilpp, Open Court, 859–1013.
- Joyce, James M. 1999. *The Foundations of Causal Decision Theory*. Cambridge University Press.
- Kahneman, Daniel, Slovic, Paul, and Tversky, Amos, eds. 1982. *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press.
- Kaplan, Mark. 1996. *Decision Theory as Philosophy*. Cambridge University Press.
- Levi, Isaac. 1986. *Hard Choices*. Cambridge University Press.
- Maher, Patrick. 1993. *Betting on Theories*. Cambridge University Press.
- . 2006a. The concept of inductive probability. *Erkenntnis* 65:185–206.
- . 2006b. Review of *Putting Logic in its Place*, by David Christensen. *Notre Dame Journal of Formal Logic* 47:133–149.

- . 2007. Explication defended. *Studia Logica* 86:331–341.
- . 2008. Physical probability. In *Logic, Methodology and Philosophy of Science: Proceedings of the Thirteenth International Congress*, eds. Clark Glymour, Wei Wang, and Dag Westerståhl, Kings College Publications. To appear.
- Ramsey, Frank P. 1926. Truth and probability. In *Studies in Subjective Probability*, eds. Henry E. Kyburg, Jr. and Howard E. Smokler, Krieger, 25–52. 2nd ed.
- Savage, Leonard J. 1954. *The Foundations of Statistics*. John Wiley. 2nd ed., Dover 1972.