

Bayesian Probability

Patrick Maher

September 4, 2008

ABSTRACT. Bayesian decision theory is here construed as explicating a particular concept of rational choice and Bayesian probability is taken to be the concept of probability used in that theory. Bayesian probability is usually identified with the agent's degrees of belief but that interpretation makes Bayesian decision theory a poor explication of the relevant concept of rational choice. A satisfactory conception of Bayesian decision theory is obtained by taking Bayesian probability to be an explicatum for inductive probability given the agent's evidence.

1 Introduction

There are at least two different senses in which a choice is ordinarily said to be rational: (1) the choice is one that could be made by an ideal agent who has the same empirical evidence that the real agent has; (2) the choice is one that the agent does not deserve blame for making. I will call a choice that is rational in sense (1) *absolutely* rational and a choice that is rational in sense (2) *deontologically* rational. The following example illustrates the difference between these concepts.¹

Andrew must decide whether to bet at even odds that a double six will be obtained in 24 tosses of a pair of dice. He is unable to calculate the probability of this in the time available but he recollects that the probability of getting a six in four tosses of a single die is greater than $1/2$ and he notices that 24 is to 36 (the number of possible outcomes of tossing two dice) as 4 is to 6 (the number of possible outcomes of tossing one die). He knows that according to an old gambling rule, it follows that the probability of obtaining a double six in 24 tosses of two dice is greater than $1/2$ and so he accepts the bet.

Andrew's choice isn't absolutely rational because an ideal agent with Andrew's evidence would have determined that the probability of getting the double six

¹ This example is loosely based on a question that was put to Pascal by the Chevalier de Méré; see Ore (1960) for an accurate historical account. In this example, and also in others that I will give later, I assume that the utility of money is a linear function of the amount of money.

is actually a little less than $1/2$. On the other hand, Andrew's choice is deontologically rational, since he did the best he could under the circumstances.

Both these concepts of rational choice are vague and unclear concepts of ordinary language. One way to theorize about a concept of this sort is to identify a clear and precise concept that is sufficiently similar to the vague concept that the former can be used in place of the latter. This methodology is called *explication*, the vague concept is called the *explicandum*, and the similar precise concept is called the *explicatum*. For further discussion of explication, see Maher (2007).

Bayesian decision theory will here be understood as proposing maximization of expected utility as an explicatum for the concept of absolutely rational choice. This is a stipulation about what I will mean by "Bayesian decision theory" in this paper. I think this stipulation is in accord with how many people have understood Bayesian decision theory but I know that some decision theorists understand the theory differently. The important point is that what I am calling "Bayesian decision theory" explicates a kind of rationality we would like our choices to have and for that reason it is worth discussing this theory.

Since expected utility is defined in terms of probability and utility, a full specification of the explicatum 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*. The question with which this paper is concerned is: What interpretation of Bayesian probability makes maximization of expected utility a good explicatum for the concept of absolutely rational choice?

2 Physical probability

A good place to begin is with the senses that the word "probability" has in ordinary language. There are two such senses, which can be distinguished by the following example: 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 the coin 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 they must involve different senses of "probability." 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 decision problems the relevant states of nature don’t have physical probabilities.² 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. Proposal 1 entails that expected utility is undefined in such decision problems. However, there will often be rational choices in such situations; for example, John may have good evidence that Mary loves him, in which case he would be rational to propose to Mary. (Here and in what follows, “rational choice” means absolutely rational choice unless otherwise indicated.)

Second, when the states do have physical probabilities, the acts that maximize expected utility using Proposal 1 are not always rational choices. For example, suppose my options are to bet at 1:2 odds that a coin will land heads, to bet at 1:2 odds that it will land tails, or not to bet, and suppose all I know about the coin is that it has either two heads or two tails. The rational choice is not to bet but, if the coin is in fact two-headed, the choice that maximizes expected utility with Proposal 1 is to bet on heads. Hence the rational choice here is not the one that maximizes expected utility.

Third, as the previous example illustrates, Proposal 1 has the result that which act maximizes expected utility is an empirical matter that may be unknown to the agent. By contrast, the act that is rational for an agent in a given situation does not depend on empirical facts about which the agent has no information.

Thus adoption of Proposal 1 would lead to a significant divergence between maximization of expected utility and rational choice, making the former not an adequate explicatum for the latter. Hence Proposal 1 is unsatisfactory.

² In decision theory, the “states of nature” are ways the world might be that, together with the act chosen, determine the consequence obtained by the agent. In what follows, I refer to them simply as “states.”

3 Inductive probability

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 . 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 sentence* for a probability concept be a sentence that says the concept has a specific value for specific arguments. For example, the following is an elementary sentence for inductive probability:

The inductive probability that a coin landed heads, given that it either landed heads or tails, is $1/2$.

I say that a probability concept is *logical* if all true elementary sentences for it are analytic.³ Every probability function whose numeric values are specified by definition is logical, so there demonstrably are probability concepts that are logical. In particular, the function c^* of Carnap (1950) is logical, since its numeric values are fixed by its definition. Similarly, every c_λ in Carnap's (1952) continuum of inductive methods is logical. On the other hand, physical probability isn't logical 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. Since the only other ordinary language concept of probability is physical probability, and it isn't logical, we can define inductive probability as the concept of probability in ordinary language that is logical.

Like many other concepts of ordinary language, the concept of inductive probability is vague. 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 $1/2$.

³ This conception of logical probability is similar to that of 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 (2006). 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 expected utility to be defined in some contexts in which physical probabilities don't exist. Another advantage is that it makes the expected utility of an act independent of empirical facts about which the agent has no information.

However, Proposal 2 does have a serious drawback. As noted above, inductive probabilities often lack precise numeric values. The definition of expected utility assumes numeric probabilities for the states, so Proposal 2 has the result that expected utility is only defined in the rather special situations in which the inductive probabilities of the states, given the agent's evidence, have precise numeric values. The concept of rational choice, by contrast, isn't restricted to those special situations. For example, the inductive probability that Mary loves John, given John's evidence, won't have a numeric value but it may still be rational for John to propose to Mary.

To deal with this drawback, one might think of representing inductive probabilities by sets of probability functions and then taking the explicatum of Bayesian decision theory to be maximization of expected utility relative to every 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.⁴

4 Explication of inductive probability

A better way to deal with the vagueness of inductive probability is to explicate this concept. I would choose as explicatum a conditional probability function p whose values are specified by definition. It follows that p is logical 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 numeric 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:

⁴ 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."

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.

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 if we have a good explication of inductive probability (and of utility) then someone who follows the recommendations of the theory will never make an irrational choice, and the clarity of the theory makes it possible to determine which acts maximize expected utility in contexts where it is otherwise unclear which acts are rational choices.

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, I was making a semantic

claim, not an epistemic or normative one, as you can verify by checking my definition of logical probability. 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 2006, 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 think that the advocates of this proposal have not properly appreciated the force of the following criticisms.

First, the agent's degrees of belief may be so vague that, in a particular decision problem, there is no fact as to which option maximizes expected utility with Proposal 4. There are nevertheless rational and irrational choices in many such situations. For example, I think most people's degree of belief in getting two sixes in 24 tosses of a pair of dice is too vague for there to be any fact as to whether accepting an even-odds bet on this maximizes expected utility with Proposal 4, though in fact that choice is irrational.

Second, real agent's degrees of belief often violate the laws of probability (Kahneman et al. 1982); following the usual terminology, I will call such degrees of belief *incoherent*. The definition of expected utility, however, assumes that the probabilities really are probabilities, that is, that they satisfy the probability calculus. (Otherwise, the expected utility of an act can depend on how one partitions the states.) Thus expected utility is undefined for such agents. Nevertheless, even for agents with incoherent degrees of belief, there are rational and irrational choices, so once again Proposal 4 makes maximization of expected utility diverge from the concept of rational choice.

Third, if the agent's degrees of belief are not suitably related to the agent's evidence, then choices that maximize expected utility with Proposal 4 can be irrational, as the following example illustrates.

Belinda has a high degree of belief that a die will land six on the next toss, for no reason except that she has a hunch this will happen, and despite the fact that most of her hunches have been incorrect in the past. On this basis, Belinda bets at even odds that the die will land six on the next toss.

Belinda's decision isn't absolutely rational because an ideal agent with her evidence wouldn't put so much weight on the hunch. It isn't even deontologically rational because Belinda can rightly be blamed for taking the hunch

so seriously, especially since she knows her poor track record with hunches. However, with Proposal 4 Belinda's decision does maximize expected utility.

These criticisms of Proposal 4 rest on facts that are mostly familiar to the many decision theorists who endorse that proposal. A common 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 incoherent; Bayesian decision theory is about what is rational and rational degrees of belief are coherent. Now this claim about what is rational may or may not be true, depending on what is meant by "rational" here. But we don't need to go into that because the response is in any case irrelevant to the objection. The objection is that, when the agent's degrees of belief are too vague, or incoherent, or not suitably related to the evidence, Proposal 4 has the result that the explicatum (maximization of expected utility) diverges from the explicandum (rational choice). Even if the agents in these situations are irrational, that doesn't change the fact that there is this divergence between explicandum and explicatum, since the concept of rational choice applies to irrational as well as rational agents. In fact, the concept would be of little use if it applied only to rational agents; that would be like having a law against murder that doesn't apply to murderers.

Some advocates of Proposal 4 respond to these kinds of criticisms by saying, in effect, that the explicandum for Bayesian decision theory is not absolutely rational choice but rather the choice that is rational given the beliefs the agent actually has. What exactly this means is unclear, but I take it that these decision theorists would say Belinda's decision is rational given her beliefs. I will call the sort of rationality exhibited by Belinda's decision *relative rationality*, since it is relative to the agent's beliefs. What these decision theorists are saying, then, is that the explicandum for Bayesian decision theory is relative rationality, not absolute rationality as I have taken it to be. My response is that this disagreement is merely a semantic one about how to use the term "Bayesian decision theory" and I settled that by stipulation in Section 1; what I mean by "Bayesian decision theory" has absolute rationality as its explicandum. If some decision theorists want to use the term "Bayesian decision theory" differently, they may do so, but that does not address the question of how best to interpret the concept of Bayesian probability for Bayesian decision theory as understood here. Furthermore, as the example of Belinda shows, relatively rational choices can be silly and even blameworthy; what we want is for our choices to be absolutely rational. Hence Bayesian decision theory as I am understanding it is a theory of the kind of rationality we care most about.

6 Rational degree of belief

It may be thought that degree of belief provides an adequate interpretation of Bayesian probability provided we add a qualification to the effect that the degrees of belief are rational. I will now explore this idea.

One might try to articulate the idea with:

Proposal 5 *Bayesian probabilities are the agent's degrees of belief, provided they are rational.*

For this proposal to be satisfactory, the term “rational” must be understood in such a way that, whenever the act that maximizes expected utility using Proposal 4 is irrational, the agent’s degrees of belief are also irrational. For example, the degrees of belief of both Andrew and Belinda must be deemed irrational. But then, in all such cases Proposal 5 has the result that expected utility is undefined. On the other hand, there are rational and irrational choices in such situations; for example, Andrew and Belinda both made irrational choices. Hence Proposal 5 has the result that the explicatum of Bayesian decision theory often fails to exist in cases where the explicandum applies, which makes this proposal unsatisfactory.

Another objection to Proposal 5 is that, even when the agent’s degrees of belief are rational in the relevant sense, they may still be vague, so that it may be unclear which acts maximize expected utility; by contrast, a good explicatum should be clear and precise.

An alternative approach is:

Proposal 6 *Bayesian probabilities are the degrees of belief that are rational for the agent to have.*

This avoids the first objection to Proposal 5 because with Proposal 6 Bayesian probabilities can exist even if the agent doesn’t have rational degrees of belief.

If Proposal 6 is to correctly handle examples like those of Andrew and Belinda, the degrees of belief that count as “rational” must be ones that are appropriately related to the agent’s evidence. But what is the appropriate relation? The best sense I can make of this is to say that a person’s degrees of belief are appropriately related to the evidence if they equal the inductive probabilities given the person’s evidence. But then Proposal 6 is just an obscure way of stating Proposal 2 and the reference to belief is otiose.

In any case, however the concept of rational degree of belief may be clarified, it is sure to be vague, so that Proposal 6 will have the consequence that the concept of expected utility lacks the clarity and precision that a good explicatum should have. A solution to this is to adopt Proposal 3, which is what I recommend.

7 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 eventually 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.
- Kahneman, Daniel, Slovic, Paul, and Tversky, Amos, eds. 1982. *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press.
- Maher, Patrick. 2006. The concept of inductive probability. *Erkenntnis* 65:185–206.
- . 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.
- Ore, Oystein. 1960. Pascal and the invention of probability theory. *American Mathematical Monthly* 67:409–419.