Problems for Bayesian Epistemology

John L. Pollock
Department of Philosophy
University of Arizona
Tucson, Arizona 85721
pollock@arizona.edu
http://www.u.arizona.edu/~pollock

In the past, few mainstream epistemologists have endorsed Bayesian epistemology, feeling that it fails to capture the complex structure of epistemic cognition. The defenders of Bayesian epistemology have tended to be probability theorists rather than epistemologists, and I have always suspected they were more attracted by its mathematical elegance than its epistemological realism. But recently Bayesian epistemology has gained a following among younger mainstream epistemologists. I think it is time to rehearse some of the simpler but still quite devastating objections to Bayesian epistemology. Most of these objections are familiar, but have never been adequately addressed by the Bayesians.

1. A Sketch of Bayesian Epistemology

Bayesian epistemology is based on two core ideas — subjective probability, and Bayesian conditionalization. I will focus on each in turn.

The theory of subjective probability begins with the common-sense idea that belief comes in degrees. Some beliefs are held more strongly, or with a greater degree of confidence or conviction, than others. This loose intuitive idea seems unproblematic. However, if it is to form the basis for a theory of numerical probabilities, it is desirable to have a way of quantifying degrees of belief. The orthodox Bayesian way of doing this is due to Frank Ramsey (1926) and Leonard Savage (1954). Their idea was that the stronger one's degree of belief in a proposition, the riskier odds one should be willing to accept when betting that the proposition is true. The orthodox Bayesian attempts to capture this by defining:

> A cognizer S has degree of belief $n/(n+r)$ in a proposition P iff S would accept any bet that P is true with odds better than r:n, and S would accept any bet that P is false with odds better than n:r.

It is crucial for their use in Bayesian epistemology that subjective probabilities conform to the probability calculus. I will adopt the weakest possible set of axioms for the probability calculus:

(1) $\text{prob}(P \,\&\, {\sim}P) = 0$;
(2) $\text{prob}(P \vee {\sim}P) = 1$;
(3) $\text{prob}(P \vee Q) = \text{prob}(P) + \text{prob}(Q) - \text{prob}(P \,\&\, Q)$.

The first two axioms normalize the values of probabilities, requiring them to fall between 0 and 1. The third axiom tells us that probability is a finitely additive measure. These axioms are often strengthened, replacing (3) with an axiom of countable additivity, and adding an axiom to the effect that logically equivalent propositions must have the same probability. But for the purposes of this paper I will adopt only this maximally abstemious set of axioms.

Bayesians sometimes speak loosely as if they identify subjective probabilities with degrees of belief. Unfortunately, it is universally agreed that real people do not have degrees of belief that conform to the probability calculus. This might be demonstrated by careful psychological experiment, as when Tversky and Kahneman (1974) showed that, given a description of a fictional woman Linda (Linda is 31 years old, single, outspoken, and very bright; she majored in philosophy; as a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations), people tend to ascribe higher probabilities (degrees of belief?) to Linda's being a feminist bank teller than to Linda's being a bank teller. That, of course, is inconsistent with the probability calculus. But there is a simpler argument demonstrating the same thing. It follows from the probability calculus that:

(3) tautologies have probability 1;
(4) tautologically inconsistent propositions have probability 0;
(5) and tautologically equivalent propositions have the same probabilities.

Thus, if degrees of belief satisfy the probability calculus, everyone would have to have a degree of belief of 1 in every tautology, a degree of belief of 0 in every tautological inconsistency, and the same degrees of belief in every pair of tautologically equivalent propositions. Recalling what this means in terms of the definition of "degree of belief", everyone would have to be willing to bet at any odds at all, no matter how disadvantageous, that any tautology is true. Furthermore, they would have to be willing to accept such bets even without knowing that the propositions *are* tautologies. It can be arbitrarily difficult to ascertain whether a complex proposition is a tautology — you needn't be able to tell at a glance — and the probability calculus assigns a tautology a probability of 1 independent of your knowing whether it is a tautology. For example, $[P \leftrightarrow (Q \& \sim P)] \rightarrow \sim Q$ is a tautology, but most people do not find this obvious without working it out. Suppose you are unaware that this is a tautology, and I give you the following proposition having this form:

> If it is true that the President is at Camp David iff both the Secretary of State is in Colorado and the President is not at Camp David, then the Secretary of State is not in Colorado.

Suppose I offer you a bet in which I pay you one cent if this proposition is true, but if it is false you allow yourself to be lowered very slowly, feet first, into a vat of sulphuric acid. Would you accept the bet? I doubt it, but then you do not have a degree of belief of 1 in this proposition.[1] Accordingly, your degrees of belief do not conform to the probability calculus.

I sometimes hear the response that this is not an interesting problem because no one knows how to handle necessary truths and tautologies in probability theory. Even if that were right, it is important to realize that this problem is not just about tautologies and necessary truths. The axioms of the probability calculus say nothing special about tautologies. All they tell us is that probability is a finitely additive measure normalized to 0 and 1. There is no way to avoid this problem without giving up finite additivity, and that is to give up the entire probability calculus.

The orthodox Bayesian response to the tautology problem is to identify subjective probabilities with the degrees of belief people "ideally ought to have" rather than their actual degrees of belief. Subjective probability becomes a limit notion, describing the limit to which a cognizer's degrees of belief approach as he performs more reasoning or engages in further epistemic cognition. Then Bayesians typically appeal to some version of the "Dutch book argument" to argue that rational cognizers *should have* degrees of belief that conform to the axioms of the probability calculus. In betting parlance, a "Dutch book" is a combination of bets on which a person will suffer a collective loss no matter what happens. For instance, suppose you are betting on a coin toss and are willing to accept odds of 1:2 that the coin will land heads and are also willing to accept odds of 1:2 that the coin will land tails. I could then place two bets with you, betting 50 cents against the coin landing heads and also betting 50 cents against the coin landing tails, with the result that no matter what happens I will have to pay you 50 cents on one bet but you will have to pay me $1 on the other. In other words, you have a guaranteed loss — Dutch book can be made against you. The Dutch book argument (due originally to Ramsey 1926) consists of a mathematical proof that if an agent's degrees of belief do not conform to the probability calculus then Dutch book can be made against him. That is, he will accept a combination of bets on which he is guaranteed to suffer a collective loss no matter what happens. It is alleged that it is irrational to put oneself in such a position, so it cannot be rational to have degrees of belief that do not conform to the probability calculus.[2]

The Dutch book argument is used to motivate the following definition of subjective probability:

> The subjective probability of *P* for a person *S* is the degree of belief *S* rationally ought to have in *P*.

The Dutch book argument purports to show that the degrees of belief a person ought to have must conform to the probability calculus.

---

[1] In fact, you would probably not accept this bet even if you believed the proposition is a tautology, so even then you would not have a degree of belief of 1.

[2] It is worth mentioning that the Dutch book argument also supports the additional axiom that logically equivalent (not just tautologically equivalent) propositions have the same probability.

Having defined subjective probability to their satisfaction, Bayesians take that to be the central concept for use in epistemology. They either identify degrees of justification with subjective probabilities or they simply dispense with talk of degrees of justification and talk instead about subjective probabilities, maintaining that the objective of an epistemic agent is to have beliefs with high subjective probabilities.

The second core idea of Bayesian epistemology is Bayesian conditionalization. This concerns how new evidence impacts a cognizer's subjective probabilities. For new evidence E, let $\text{prob}_E(P)$ be the subjective probability of P given the evidence E. The principle of Bayesian conditionalization then proposes that $\text{prob}_E(P) = \text{prob}(P/E)$. One way to think of this is that acquiring the evidence E has the consequence that $\text{prob}_E(E) = 1$. If acquiring the evidence leaves the conditional probability $\text{prob}(P/E)$ unchanged, then by the probability calculus

$$\text{prob}_E(P) = \text{prob}_E(P/E) \cdot \text{prob}_E(E) + \text{prob}_E(P/{\sim}E) \cdot (1 - \text{prob}_E(E)) = \text{prob}_E(P/E) = \text{prob}(P/E).$$

So the substantive assumptions underlying Bayesian conditionalization are that (a) evidence has probability 1, and (b) the conditional probability $\text{prob}(P/E)$ does not change when we acquire the new evidence E.

Bayesian conditionalization should not be confused with Bayes' Theorem. Bayes' Theorem is simply a theorem of the probability calculus, saying:

$$\text{prob}(P / E) = \text{prob}(E / P) \cdot \frac{\text{prob}(P)}{\text{prob}(E)}.$$

This can be coupled with Bayesian conditionalization to conclude that

$$\text{prob}_E(P) = \text{prob}(E/P) \cdot \frac{\text{prob}(P)}{\text{prob}(E)}.$$

$\text{prob}(P)$ and $\text{prob}(E)$ are the "prior probabilities" of P and E, i.e., their probabilities prior to acquiring the new evidence, and $\text{prob}(E/P)$ is the prior probability of acquiring the evidence if P is true. If the latter probabilities can be computed, this gives us a way of computing $\text{prob}_E(P)$.

The combination of Bayes' theorem and Bayesian conditionalization has played an important role in the philosophy of science. Among other things it is often argued, by appealing to a theorem of de Finetti (1974), that Bayesian conditionalization provides a solution to the problem of induction.

Because subjective probabilities satisfy the probability calculus, Bayesians can prove interesting and often surprising theorems about subjective probabilities that appear to have important implications for epistemology.[3] The mathematical power of Bayesian epistemology makes it a very appealing tool, but as I will now show, there are major, and I think insurmountable, problems for the foundations of Bayesian epistemology.

2. Degrees of Belief

Recall that the official definition of "degree of belief" is:

> A cognizer S has degree of belief $n/(n+r)$ in a proposition P iff S would accept any bet that P is true with odds better than r:n, and S would accept any bet that P is false with odds better than n:r.

This attempts to assign numerical values to degrees of belief in terms of what bets an epistemic cognizer would be willing to accept. r and n are the *stakes* and the ratio r:n is the *odds*. This definition requires that the acceptability of a bet should be a function only of the odds, and not of the absolute values of the stakes. The problem is that well known empirical results in economics demonstrate that real people do not have degrees of belief in this technical sense. As economists well know, when you vary the stakes of a bet but keep their ratio (the odds) fixed, people may change their willingness to bet. Suppose I ask you where you left your car keys, and you reply, "On my

[3] For example, see Skyrms (1980), Howson and Urbach (1989), and White (2006).

dresser". If you are sufficiently confident of this, and I offer you a bet in which I pay you ten cents if you are right but you pay me one dollar if you are wrong, you may be willing to accept this bet. But suppose I offer you a bet in which I pay you ten thousand dollars if you are right but you pay me one hundred thousand dollars if you are wrong. You may not be willing to accept this bet, and yet the odds have not changed. Only the stakes have changed.

The upshot is that real people do not have degrees of belief in the orthodox Bayesian sense. However, subjective probability is defined in terms of degrees of belief. It is supposed to be the limit to which degree of belief goes as the cognizer does more of the relevant reasoning. It is not clear whether this is a fatal problem for orthodox Bayesianism, because subjective probability is not identified with actual degree of belief, but rather with the degree of belief the cognizer rationally ought to have. It could be argued that although real people do not have degrees of belief, they *should* have degrees of belief, i.e., the betting behavior should not be affected by the stakes. This claim cannot be defended by appealing to the Dutch book argument, the Bayesian theory becomes as problematic as the alternatives.

3. Degrees of Belief and the Probability Calculus

One of the main attractions of the orthodox notion of degree of belief is that, if people actually had degrees of belief, we could measure them experimentally just by examining their betting behavior. This provides an elegant solution to the problem of assigning numerical values to degrees of belief. The orthodox Bayesian then goes on to define subjective probabilities to be the degrees of belief that a cognizer rationally ought to have. Historically, what attracted people to subjective probabilities in the first place was that they seemed to provide an elegant way of defining probability in terms that no one found problematic, whereas competing objective theories always appealed to concepts that seemed as problematic as the probability concept being defined. However, if people do not have degrees of belief in the technical Bayesian sense, that seems to be a serious problem for this definition.

On the other hand, it must be acknowledged that people do have degrees of belief in a looser more informal sense (degrees of confidence, or degrees of conviction). Perhaps we can reformulate Bayesian epistemology in terms of one of these more informal notions. It becomes problematic just how we can assign numerical values to the more informal notions, but for the moment let us waive that difficulty and see how well the rest of the Bayesian theory works if we suppose we have numerical values for the degrees of belief of real people.

If people's degrees of belief (however they are defined) conformed to the probability calculus, we could simply identify subjective probabilities with degrees of belief, and when they are being careless, Bayesians occasionally talk as if that is what they are doing. Unfortunately, as in section one, even if there were no problems regarding the orthodox definition of degrees of belief, real people's degrees of belief would not conform to the probability calculus. On the orthodox construal of degrees of belief, that would require being willing to bet at any odds at all, no matter how disadvantageous, that any tautology is true. Furthermore, we would have to be willing to accept such bets even without knowing that the propositions *are* tautologies. Employing a more informal notion of degree of belief does not seem to help. If such a notion (however numerical values are to be assigned to it) conformed to the probability calculus, you would have to have absolute conviction or total confidence in the truth of tautological propositions even without knowing that they are tautologies. But no one does. I will refer to this as *the tautology problem*.

The result is that if we simply identify subjective probabilities with degrees of belief, then on any reasonable sense of degree of belief, subjective probabilities will not satisfy the probability calculus. It is worth rehearsing why that is a problem. One of the main appeals of Bayesian epistemology is the mathematical power the use of probabilities gives us. Given the probability calculus and Bayesian conditionalization, Bayesians purport to prove many interesting and powerful theorems about epistemological concepts like theory confirmation. If we forsake the probability calculus, we give up the main appeal of Bayesian epistemology.

Bayesians are well aware of this, and so (at least when they are being careful), they do not identify subjective probabilities with the degrees of belief of real cognizers. Instead, as we have noted, the orthodox Bayesian strategy is to identify subjective probabilities with the degrees of belief a cognizer rationally should have in propositions rather than the degrees of belief he actually has.

4. From Degrees of Belief to Subjective Probabilities

It is not wholly implausible to suppose that however we ultimately decide to define "degree of belief", there will be empirical techniques enabling us to measure the degrees of belief real cognizers have in many propositions. So let us simply grant that for the sake of argument. Subjective probabilities are then taken to be the degrees of belief the cognizer rationally should have in propositions. But what determines what degrees of belief cognizers should have? No Bayesian has ever proposed an answer to this question. Bayesians typically propose what they regard as rational constraints on degrees of belief, the claim being that agents whose degrees of belief do not satisfy these constraints are irrational. The weakest constraint, that all Bayesians endorse, is that degrees of belief should satisfy the probability calculus, and they appeal to the Dutch book argument to justify this claim. But given that the degrees of belief of a real cognizer will not satisfy the probability calculus, how should we repair them so that they do? No one has ever proposed a unique way of doing this. For any contingent proposition P and any number r such that $0 < r < 1$, there will be a way of modifying a cognizer's actual degrees of belief so that they satisfy the probability calculus and the degree of belief in P is r. To illustrate, suppose the cognizer only has degrees of belief for four propositions: $\text{db}(P) = .5$, $\text{db}(Q) = .7$, $\text{db}(P \vee Q) = .8$, and $\text{db}(P \,\&\, Q) = .3$. These degrees of belief do not conform to the probability calculus, because

$$\text{db}(P \vee Q) \neq \text{db}(P) + \text{db}(Q) - \text{db}(P \,\&\, Q).$$

If our only requirement is that the degrees of belief be altered to conform to the probability calculus, then there are infinitely many ways of accomplishing this. For example, we could set

$$\text{db}(P) = .7,\ \text{db}(Q) = .8,\ \text{db}(P \vee Q) = .9,\ \text{and}\ \text{db}(P \,\&\, Q) = .6.$$

Even if we require that the number of changes be minimal, there are four ways of accomplishing this:

$$\text{db}(P) = .4,\ \text{db}(Q) = .7,\ \text{db}(P \vee Q) = .8,\ \text{and}\ \text{db}(P \,\&\, Q) = .3.$$

$$\text{db}(P) = .5,\ \text{db}(Q) = .6,\ \text{db}(P \vee Q) = .8,\ \text{and}\ \text{db}(P \,\&\, Q) = .3.$$

$$\text{db}(P) = .5,\ \text{db}(Q) = .7,\ \text{db}(P \vee Q) = .9,\ \text{and}\ \text{db}(P \,\&\, Q) = .3.$$

$$\text{db}(P) = .5,\ \text{db}(Q) = .7,\ \text{db}(P \vee Q) = .8,\ \text{and}\ \text{db}(P \,\&\, Q) = .4.$$

These all conform to the probability calculus. Which is the uniquely rational set of degrees of belief that the cognizer ought to have? If the only constraint on degrees of belief is that they satisfy the probability calculus, then there is no uniquely rational set of degrees of belief. But then subjective probabilities do not exist for cognizers whose degrees of belief do not initially conform to the probability calculus. I will refer to this as *the uniqueness problem*.

Could we avoid the uniqueness problem by adding some more constraints on degrees of belief? A few additional constraints have been proposed. Bayesian conditionalization is one, but it only yields degrees of belief that conform to the probability calculus if it starts with degrees of belief that conform to the probability calculus. David Lewis (1981, 1994) proposed what he called The Principal Principle, and Bas van Fraassen (1981) proposed his Reflection Principle, but none of these are any help in producing a unique set of degrees of belief. Nor were they intended to be. Bayesians have simply ignored the uniqueness problem. But it is a crippling problem for Bayesian epistemology. Without some unique way of rendering a cognizer's degrees of belief consistent with the probability calculus, there is no such thing as *the* degree of belief the cognizer ought to have in a proposition. In other words, there is no such thing as the subjective probability of a proposition for a cognizer whose degrees of belief do not conform to the probability calculus (which is all of us).

5. Ideal Cognizers

Some Bayesians (e.g., Skyrms 1980, 1984; and Lewis 1981) try to avoid the tautology problem and the uniqueness problem by insisting that they are only talking about ideal cognizers. Ideal cognizers have no constraints on their reasoning powers, memory, etc. The main thing this is supposed to achieve is that because ideal cognizers can instantly recognize when their degrees of belief fail to conform to the probability calculus, and they know (via the Dutch book argument) that their degrees of belief should conform to the probability calculus, they will never have degrees of belief that fail to satisfy the probability calculus. These Bayesians propose that for ideal cognizers, we can simply identify subjective probabilities with actual degrees of belief.

Even if this were true, what does it have to do with us — real cognizers with finite reasoning powers, limited memory capacity, etc? For instance, if a theory of ideal agents says that they should attend to all of the logical consequences of their beliefs, but we as human beings cannot do that, then the recommendations applicable to ideal agents are simply not applicable to us. *We* should do something else. This point can be illustrated by reconsidering the applicability of the Dutch book argument to real cognizers. It is certainly undesirable to have Dutch book made against you, but is it truly irrational to have degrees of belief making that possible? Real cognizers, with limited cognitive powers, will not be able to avoid having degrees of belief that violate the probability calculus. For instance, if P and Q are tautologically equivalent propositions, but I am unaware of this and have different degrees of belief in P and Q, Dutch book can be made against me. But it is not as if I am making an egregious mistake in reasoning. I am just subject to unavoidable ignorance. No real cognizer can have knowledge of every tautological equivalence. That is computationally impossible. Rationality might prescribe *trying to find out* whether P and Q are tautologically equivalent, but that can take arbitrarily much reasoning, and it may be impossible to complete the reasoning before it is time to bet. So one cannot be rationally culpable for failing to discover the equivalence.

As I use the term "theory of rational belief", it is about what we, and other resource-bounded cognizers, should do. I want to know how, given our cognitive limitations, we should decide what beliefs to hold. In other words, I want a theory of *real rationality* as opposed to a theory of *ideal rationality*. This distinction is widely recognized, but it often seems to be supposed that as philosophers our interest should be in ideal rationality. The rationality a human can achieve is mere "bounded rationality" — a crude approximation to ideal rationality (Cherniak 1986). But surely we come to epistemology with an initial interest in what we, and cognizers like us, should believe. This is the notion of rationality that first interests us, and this is what I am calling "real rationality".

Although theories of ideal cognizers are not directly about how real cognizers should perform cognitive tasks, a plausible suggestion is that the rules of rationality for real cognizers should be such that, as we increase the reasoning power of a real cognizer, insofar as he behaves rationally his behavior will approach that of an ideal rational cognizer in the limit. This is to take theories of ideal rationality as imposing a constraint on theories of real rationality. However, even if this is right, it does not seem to be of any help in characterizing subjective probabilities for real resource-bounded cognizers. The difficulty is that if the only rational constraints on a cognizer's degrees of belief are the familiar ones Bayesians have discussed, then there is more than one way of altering a real cognizer's degrees of belief to turn him into an ideal cognizer. Thus there is no way to pick a set of ideal degrees of belief as the unique set of degrees of belief that would be possessed by the cognizer if he did enough reasoning to render his degrees of belief consistent with the probability calculus and hence become an ideal cognizer.

Thus far my conclusion is that orthodox Bayesian epistemology can only make sense of subjective probabilities for ideal cognizers, but what we need for epistemology (viewed as part of a theory of real rationality) is subjective probabilities for real resource-bounded cognizers. The orthodox Bayesian theory provides no way of making sense of the latter.

6. Subjective Probabilities from Preference Rankings

Thus far I have focused on what I called "the orthodox version" of Bayesian epistemology, which attempts to define subjective probabilities in terms of degrees of belief. There is, however, another approach that has been at least as influential and avoids appeal to degrees of belief and the Dutch book argument. What this approach takes as primitive is the cognizer's preferences. It is assumed that cognizers have preferences not just between the outcomes of bets, but between bets themselves. For example, a cognizer might reasonably prefer a bet in which he receives ten dollars if a coin lands heads but loses five dollars if it lands tails to one in which he receives five dollars if the coin lands heads and loses five dollars if it lands tails.

The orthodox definition of "degree of belief" assumes that we can assign numbers to the values of the outcomes of bets. Economists and decision-theorists often regard that assumption as problematic, on the grounds that we cannot introspect numerical values. It is often maintained that all we can introspect are binary preferences.[4] For each cognizer S, the relation "S prefers x to y" is a binary non-numerical relation. If we consider the algebraic properties of this relation, it is plausible that certain properties are required by rationality. For example, it is plausible to suppose that rationality requires our preference relations to be transitive — if S prefers x to y and prefers y to z, then S ought to prefer x to z. It is also plausible that rationality imposes some constraints specifically pertaining to preferences between bets. A plausible example might be:

> If S prefers x to y, then for any z, S should prefer a bet in which he receives x if P is true and pays out z if P is false to a bet in which he receives y if P is true and pays out z if P is false.

Ramsey (1926) was the first to realize that it might be possible to find a set of constraints on preferences such that (1) a cognizer is irrational insofar as his preferences violate the constraints, and (2) if a cognizer's preferences satisfy the constraints, it is possible to recover numerical values for both the cognizer's subjective probabilities and his valuations of outcomes.

The mechanism for recovering numerical values for probabilities and valuations assumes the standard *optimality prescription* of decision theory. We assume that our task is to choose an action from a set $\mathbf{A} = \{A_1,\ldots,A_n\}$ of *alternative actions*. The actions are to be evaluated in terms of their outcomes. We assume that the *possible outcomes* of performing these actions are partitioned into a set $\mathbf{O} = \{O_1,\ldots,O_m\}$ of pairwise exclusive and jointly exhaustive outcomes. In other words, it is logically impossible for two different members of $\mathbf{O}$ to both occur, and it is logically necessary that some member of $\mathbf{O}$ will occur. We further assume that we know the probability $\mathrm{prob}(O/A)$ of each outcome conditional on the performance of each action. Finally, we assume a *utility measure* $\mathrm{U}(O)$ assigning a numerical value to each possible outcome. The *expected value* of an action is defined to be a weighted average of the values of the outcomes, discounting each by the probability of that outcome occurring if the action is performed:

$$\mathrm{EV}(A) = \mathrm{U}(O_1)\cdot\mathrm{prob}(O_1/A) + \ldots + \mathrm{U}(O_m)\cdot\mathrm{prob}(O_m/A).$$

According to the optimality prescription, actions are to be compared in terms of their expected values, and rationality dictates choosing an action that is *optimal*, that is, one such that no alternative has a higher expected value.[5]

Ramsey's idea was to find a set of constraints such that (1) a cognizer is irrational insofar as his preferences violate the constraints, and (2) if a cognizer's preferences satisfy the constraints, then there is a function prob assigning real-numbers to propositions and a function U assigning real numbers to the items in the cognizer's preference ranking such that (a) prob satisfies the probability calculus, (b) the cognizer prefers x to y iff $\mathrm{U}(x) > \mathrm{U}(y)$, and (c) the value U assigns to a bet is the expected value of making the bet (computed as above in terms of prob and U). If we assume that actions are chosen on the basis of being preferred to their alternatives, and for any action A, $\mathrm{U}(A)$ is the expected-value of A defined as above, then it follows that the cognizer makes his decisions in accordance with the optimality prescription. Ramsey argued further that for the rationality constraints that he proposed, the values of prob are uniquely determined. Although the values for U are not uniquely determined, different choices of U differ only by a linear change in "scale". That is, if U_1 and U_2 are two such assignments, there are real-numbers a and b such that for any x, $\mathrm{U}_2(x) = a\cdot\mathrm{U}_1(x) + b$.

A theorem like the above, relating rationality constraints on preferences to numerical values for prob and U is called a *representation theorem*. Ramsey proved the first representation theorem, but there are disagreements about what the proper rationality constraints should be. In response to these disagreements, a number of more recent authors have proven other representation theorems for different sets of constraints (Savage 1954, Bolker 1966, Luce and Krantz 1971, Jeffrey 1983, Joyce 1998). The attraction of this general approach is that in defining prob (subjective probability) it makes no appeal to a cognizer's degrees of belief. The rational constraints are just constraints on preferences, and it is assumed that this notion is unproblematic.

The difficulty with the appeal to representation theorems is that they prove too much.[6] Because this is required by the probability calculus, any representation theorem must have the consequence

[4] But see Pollock (2006), 24-35.

[5] Although the optimality prescription is generally assumed in these discussions, Pollock (2006) proposes a number of counter-examples and suggests an alternative form of "plan-based" decision-theoretic reasoning.

that if P and Q are tautologically equivalent, then prob(P) = prob(Q), and hence the expected-value of a bet that turns upon whether P is true will be the same as the expected-value of an otherwise identical bet that turns upon whether Q is true. It follows that if the cognizer's preferences satisfy the rationality constraints, then he will be indifferent between the two bets. However, if the cognizer does not realize that P and Q are tautologically equivalent then he may regard one of P and Q as more probable than the other and hence not be indifferent between the bets. For familiar reasons, cognizers will be ignorant of most tautological equivalences, and so will almost certainly have preferences that do not satisfy the constraints. This is just the tautology problem again. Appeal to preference rankings does not avoid it.

The upshot is that representation theorems enable us to define subjective probability for cognizers whose preferences satisfy the constraints, but only ideal cognizers can be expected to have preferences satisfying the constraints. So once again, this approach might work for defining subjective probabilities for ideal cognizers, but it is of no help in defining subjective probability for real resource-bounded cognizers unless we have some way of saying how real cognizers should modify their preference rankings to make them consistent with the constraints. Unless there is a unique way of doing this, appeal to ideal cognizers does not give us a way of defining subjective probabilities for real cognizers. This is the uniqueness problem again.

7. Subjective Probabilities as Degrees of Justification

Thus far I have considered problems for the standard approaches to defining subjective probability, either in terms of degrees of belief or in terms of rational constraints on preference rankings. There is a final approach to consider. A few philosophers (van Fraassen 1985, Williamson 2000, perhaps Lewis 1994) take the probability calculus to be the logic of degrees of justification, and so identify subjective probability (they often use the term "credence" instead) with the subject's degree of justification. Defining the notion "degree of justification" is itself a difficult philosophical task, but a Bayesian epistemologist might reason that any epistemologist must grant that this concept makes sense and so be willing to countenance its use in Bayesian epistemology.

However, we can resurrect the now familiar tautology problem to show that degrees of justification do not satisfy the probability calculus. If degrees of justification did satisfy the probability calculus, then any cognizer would automatically be fully justified in believing any tautology, whether or not he had reason to believe it to be a tautology. That is clearly incorrect. At least for most tautologies, we only become justified in believing them by giving an argument that shows they are tautologies.

What a cognitive agent is currently justified in believing depends both on what information he has acquired from perception (construed broadly) and how much reasoning he has performed. An agent can never perform all potentially relevant reasoning, because there is always infinitely much reasoning waiting to be performed. As the agent performs more reasoning, he may justify defeaters for earlier reasoning, with the result that previously justified beliefs become unjustified even without any new perceptual input. Still further reasoning might defeat the defeating arguments, reinstating the original beliefs, and so on. Thus as reasoning proceeds, the justifiedness of a belief may fluctuate indefinitely many times. We can, however, consider what would happen if the cognizer *could* complete all the relevant reasoning. Let us say that a proposition is currently *warranted* for a cognizer iff it would be justified if the cognizer did complete all possible relevant reasoning (without acquiring any new perceptual input). For any epistemological theory that gives reasoning pride of place, the notion of warrant will be well-defined.[7]

This suggests a strategy for the Bayesian. Rather than identifying subjective probabilities with degrees of justification, they might be identified with degrees of warrant. This automatically avoids the tautology problem, because if a cognizer does all possible relevant reasoning, he will know whether a proposition is a tautology. Note that appealing to warrant is, in effect, to define subjective probability in terms of ideal agents, because the warranted propositions are just those ideal agents would be justified in believing. However, unlike previous attempts to define subjective

[6] Another respect in which representation theorems prove too much is that they entail the optimality prescription, but as argued in Pollock (2006), the optimality prescription is only correct for simple cases. Complex cases require a form of decision-theoretic reasoning that looks are plans rather than isolated actions. It is argued, further, that decision-theoretic reasoning cannot be viewed as a search for optimal plans. Thus the rationality constraints employed in existing representation theorems cannot be correct.

[7] This notion of warrant comes from Pollock (1986). It should not be confused with the subsequently defined use of "warrant" by Plantinga (1993).

probabilities for real agents in terms of corresponding ideal agents, this strategy avoids the uniqueness problem. It does so by simply assuming that we have a prior characterization of warrant, provided by some independent epistemological theory, and then identifies subjective probability with degree of warrant. Note, however, that although this strategy avoids both the tautology and uniqueness problems, it does so by importing a great deal of prior epistemology into the theory, so the Bayesian can no longer claim to be replacing traditional approaches to epistemology with an appeal to subjective probabilities.

This approach avoids the ill-foundedness of the previous attempts to define subjective probability for real agents in terms of ideal agents, but does it give us a workable notion of subjective probability? In particular, is there any reason to think that degree of warrant satisfies the probability calculus and hence is a reasonable candidate for identification with subjective probability? Degree of warrant is defined in terms of degree of justification. It is the limit to which degree of justification goes as reasoning proceeds. But we know that degree of justification does not satisfy the probability calculus, so why should we think that degree of warrant will?

Might we employ some version of the Dutch book argument here? (David Lewis (1994) seemed to have in mind to do that.) The standard Dutch book argument defines subjective probabilities in terms of degrees of belief, which are defined in turn in terms of the agent's betting behavior. That automatically has the effect that an ideal agent will accept a bet iff the expected-value of accepting the bet, defined in terms of subjective probabilities, is higher than the expected-value of not accepting the bet. That is what makes the argument work. To make the argument work when we instead identify subjective probabilities with degrees of warrant, we need the assumption that an ideal agent would accept a bet iff the expected-value of accepting the bet, defined in terms of degrees of warrant, is higher than the expected-value of not accepting the bet. But that assumption is only reasonable if degrees of warrant can be identified with the probabilities an ideal agent would employ in decision-theoretic reasoning. Whether degrees of warrant work like probabilities at all is precisely what is at issue, and that is what the Dutch book argument is supposed to establish. Consequently, appeal to a Dutch book argument here would be circular.

It is controversial whether degrees of warrant satisfy the probability calculus. The debate about this turns largely on issues first broached by Henry Kyburg (1970). If prob(P) and prob(Q) are less than 1, it follows from the probability calculus that prob(P&Q) will normally be less than either prob(P) or prob(Q).[8] Kyburg was not a Bayesian, but he did believe that (what in his theory played the role of) degrees of warrant were probabilities, and so he bit the bullet and insisted that even if we are warranted in believing both P and Q, it does not follow that we are warranted in believing (P&Q). According to Kyburg, inferring (P&Q) from P and Q is a logical fallacy, which he dubbed "conjunctivitis".

Kyburg's claims have inspired much disagreement within epistemology. Many epistemologists take it to constitute a reductio ad absurdum of the view that degree of warrant satisfies the probability calculus. After all, the inference from P and Q to (P&Q) is just the classical rule of adjunction (or &-introduction). How could any inference be more justified than that? But other epistemologists have sided with Kyburg, insisting that adjunction is not a valid inference rule.

Let us say that an inference scheme $P_1,\ldots,P_n \vdash Q$ is *probabilistically valid* iff it follows from the probability calculus that prob(Q) is at least as great as the minimum of prob(P_1),…,prob(P_n). Kyburg's objection to adjunction amounts to observing that it is not probabilistically valid. Philosophers who have been tempted to agree with Kyburg have generally overlooked the fact that most of our standard inference rules are probabilistically invalid. For example, *modus ponens* and *modus tollens* are probabilistically invalid. In fact, no deductive inference rule having multiple premises essentially (i.e., the rule is no longer deductively valid if any premise is deleted) can be probabilistically valid. This is because the probability calculus can only guarantee that the conclusion is as probable as the *conjunction* of the premises, so problems for adjunction affect all of these other rules as well. The upshot is that if we follow Kyburg, deductive reasoning can play little role in justifying beliefs. Non-deductive (defeasible) inference rules are not probabilistically valid either, so reasoning is robbed of its role in epistemology.

This is at least odd. We tend to think of deductive reasoning as the most secure kind of reasoning. But for the kind of Bayesian we are now considering, it is more than odd — it is catastrophic. This is because the definition of warrant takes for granted the role of reasoning in the evolution of the set of justified beliefs. If reasoning has no such role, then warrant makes no sense, and consequently we cannot identify subjective probability with degree of warrant.

[8] More precisely, if prob(P/Q) < 1 and prob(Q/P) < 1 then prob(P&Q) < prob(P) and prob(P&Q) < prob(Q).

We are left with a dilemma. If reasoning does not play its normally assumed role in epistemology, then we cannot define warrant as we did, and so we cannot identify subjective probability with degree of warrant. On the other hand, if reasoning does have its normally assumed role, then inference rules like modus ponens and adjunction must convey warrant from their premises to their conclusions, and as a result degree of warrant cannot satisfy the probability calculus, in which case we again cannot identify subjective probability with degree of warrant. Thus on either horn of the dilemma, the attempt to ground Bayesian epistemology in some independent account of warrant is doomed to failure.

There remains one possibility for slipping between the horns of this dilemma. Kyburg assumed that if an inference does not preserve degree of warrant, then it should be deemed invalid and we should not employ it in reasoning. However, we might resist this assumption. For an agent that reasons defeasibly, it can be important to continue reasoning from a conclusion even though (given the reasoning done so far) the conclusion is defeated. To illustrate, suppose we construct the two arguments (on the left and on the right) in figure 1. Here the thin arrows represent inferences and the thicker arrows represent defeat relations. This is a case of "collective defeat", where each argument defeats the other and so the conclusions of both arguments should be deemed defeated. Suppose we construct the leftmost argument first, and then begin constructing the rightmost argument. When we get to *D*, we observe that it is defeated by the first argument. If we stop there we will never get to *~B* and discover that the first argument is also defeated. So when reasoning defeasibly, we cannot stop reasoning from a conclusion just because our current reasoning renders it defeated. Hence argument construction and the computation of defeat status are two separable cognitive operations.

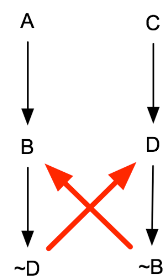

Figure 1. Collective defeat

Next notice that the construction of an argument that supports defeaters only weakly may lower the degree of justification of the conclusion without defeating it outright. Consequently, the computation of defeat status must, in the general case, be regarded as part of the computation of degrees of justification. Once we have separated argument construction from the computation of degrees of justification, a possible view would locate the problems stemming from conjunctivitis as pertaining to the computation of degrees of justification rather than argument construction. "Probabilistically invalid" inferences might still be good steps to include in arguments because, although they produce conclusions that are not as well justified as the premises, the conclusions might still be sufficiently justified to be useful. On this view, reasoning can look just the way we normally suppose, but the degrees of justification of the conclusions may decrease systematically as the reasoning progresses.

The upshot is that, contra Kyburg, it is in principle possible to retain a traditional view of argument construction while agreeing that the degrees of justification of conclusions may be less than the degrees of justification of the premises, and as a result, degrees of warrant might conceivably work like probabilities.

Although this is a technical possibility, it remains to be seen whether degrees of warrant actually do work this way. There are arguments that convince me that they do not. Consider simple counting tasks. Suppose you are counting apples in a barrel. There are, in fact, 100 apples in the barrel, and you remove each in turn, examine it to see that it is an apple, and then add it to your count. How many apples are you warranted in believing were in the barrel? Let us suppose that the probability of losing count is so low that it can be ignored. The only source of uncertainty concerns whether something is an apple. Let us suppose you are warranted in believing of each apple that it is an apple. Presumably, for the Bayesian this means that there is a probability p such that your subjective probability of its being an apple is at least p. Similarly, you are warranted in believing that there are at least r apples in the barrel iff your subjective probability of there being at least r apples is at least p. Define:

$$\text{prob-r}(p,r,N) = \sum_{i=r}^{i=N} p^i(1-p)^{n-i}\frac{N!}{i!(n-i)!}$$

By the probability calculus, the probability of there being at least r apples in the barrel is prob-r(p,r,100). Let prob-max(p,N) = the largest $r \leq N$ such that prob-r(p,r,N) $\geq p$. Then if the probability required for warrant is p, prob-max(p,100) is the largest r such that you are warranted in believing there are at least r applies in the barrel. What should the value of p be? Presumably it is .9 or higher. Computing prob-max(p,100) for some choices of p:

prob-max(0.9,100) = 86
prob-max(0.95,100) = 91
prob-max(0.99,100) = 96
prob-max(0.999,100) = 98
prob-max(0.9999,100) = 99
prob-max(0.99999,100) = 99
prob-max(0.999999,100) = 99
prob-max(0.9999999,100) = 99

So no matter how high the probability threshold required for warrant, you are never warranted in thinking there are 100 apples in the barrel, and for probabilities no higher than than .999 (which is already very high), you are not even warranted in thinking there are at least 99 apples in the barrel. In other words, it is impossible to count the number of apples in a barrel containing 100 apples.

If this does not seem sufficiently absurd, consider an even simpler counting problem. Suppose you have three daughters, Anne, Stephanie, and Lisa. You know (and are warranted in believing) of each that she is your daughter. How many daughters are you warranted in believing you have? According to the Bayesian, the answer is prob-max(p,3). Computing prob-max(p,3) for several choices of p;

prob-max(0.9,3) = 2
prob-max(0.95,3) = 2
prob-max(0.99,3) = 2
prob-max(0.999,3) = 2
prob-max(0.9999,3) = 2
prob-max(0.99999,3) = 2
prob-max(0.999999,3) = 2
prob-max(0.9999999,3) = 2

So you are only warranted in believing you have at least 2 daughters. Imagine Stephanie telling her friends, "My dad is a crazy philosopher. He knows each of us to be his daughter, but he claims that as far as he knows he only has two daughters."

These consequences of identifying subjective probability with degree of warrant are extremely counter-intuitive. They suggest very strongly that degrees of warrant do not diminish in the way the probability calculus would require.

8. Beliefs and Probabilities

Arguments like the preceding, which appeal to our intuitive judgments regarding what we are justified in believing, seem to show that neither degrees of justification nor degrees of warrant satisfy the probability calculus. In particular, if you are warranted in believing P and Q, it seems that you should also be warranted in believing (P&Q). However, there is a counter-argument that many have found persuasive. We begin with the observation that we regularly employ some kind of probability in decision-theoretic reasoning. Bayesians claim that these are subjective probabilities, but even if it is agreed that subjective probabilities make no sense for real cognizers, there must be some other kind of probability that does. Various theories have been advanced regarding non-

Bayesian probability. I will say a bit more about this below, but for now let us just suppose we have an appropriate kind of probability prob.

Here is the problem. Regardless of whether degree of warrant works like a probability, prob does. So if we begin with a number of intuitively independent propositions each having probability less than 1 and conjoin them, the probability of the conjunction should be less than the probability of the conjuncts. By making the conjunction long enough, we can make the probability of the conjunction arbitrarily small. But surely, if we know that something is very improbable, we cannot be warranted in believing it. So this seems to show that the degree of warrant of a conjunction should also decrease as we add conjuncts. And this flies in the face of the earlier intuitive arguments to the contrary. One does not have to be a Bayesian to find this argument persuasive. This argument can be run using any kind of probability.

How can we resolve this conflict? It is implausible to insist that we can remain warranted in believing something even when we know that its probability is arbitrarily low. But there is one other way out of the argument. If warranted propositions have probability 1, then so do conjunctions of warranted propositions, so contrary to supposition, conjoining warranted propositions will not produce improbable conjunctions.

Can it reasonably be maintained that warranted propositions have probability 1? I think it can. The first thing to notice is that the probabilities to be used in decision-theoretic reasoning must be sensitive to the agent's epistemic state. What an agent should do is a function in part of what he knows or believes about his current situation. Suppose you are deciding whether to carry an umbrella today. If you check the weather forecast in the morning paper and it says it will not rain, you can rationally decide to leave the umbrella at home. But then if you look out the window and see big black thunderclouds rolling towards you from the horizon, you will rationally change our mind and take the umbrella. This illustrates that the probabilities of use in decision-theoretic reasoning must be at least partly epistemic. They cannot be purely objective probabilities.[9]

The Bayesian tried to make decision-theoretic probabilities entirely epistemic, with no objective ingredient, but as we have seen, that did not work. An alternative is to employ a kind of probability that is partly objective and partly epistemic. To see how this might work, notice that beliefs and probabilities play different roles in decision-theoretic reasoning. Beliefs frame the problem, providing the background against which we compute probabilities and expected-values and make our decisions.[10] For example, suppose you are in San Francisco and you about to drive over the Golden Gate Bridge in order to visit Pt. Reyes National Seashore. In making the decision to take this little vacation trip, you consider how much it is apt to cost, how much time it is apt to take, how much pleasure you expect to derive from it, and so forth. These are all assigned probabilities. But you take it for granted that the bridge will not collapse while you are driving over it. You do not assign a probability to that. If you did assign a probability to the bridge collapsing, and it were greater than 0, then no rational person would make the trip, because no matter how improbable a bridge collapse is, dying because you are on a collapsing bridge more than outweighs whatever pleasure you may get from visiting the sea shore, and so the expected-value of making the trip would be negative. It can only be reasonable to make such a trip if you take the probability of a bridge collapse to be 0. Similarly, when you fly to a conference, you simply believe that the airplane will not crash. Unless you assigned that probability 1, the expected-value of going to the conference would be negative, no matter how good the conference. All decision problems are framed against a background of assumptions that you simply believe outright, and you take yourself to be fully justified in believing them. You do not assign probabilities to these beliefs when computing expected-values, or if you do, you assign them probability 1.

This may seem puzzling, because after all, bridges do occasionally collapse and airplanes do occasionally crash. But you regard yourself as fully justified in believing that *this bridge* will not collapse while you are on it, and *this airplane* will not crash during your flight. How can this be? This reflects a distinction between two kinds of probabilities. For decision-theoretic reasoning, we need *singular probabilities*, which attach to individual propositions and tell us how likely they are to be true. But many kinds of statistical/inductive investigations apprise us instead of *generic probabilities*, which relate properties and relations rather than propositions.[11] For example, medical research may inform us that the probability of a person with symptoms S having disease D is .6. This is not about any particular person with symptoms S — it is a general truth about the property of being a person

[9] See Pollock (2006), chapter six, for a more extensive discussion of this point.
[10] See Williamson (2000) for a similar observation.
[11] In earlier publications (Pollock 1984, 1990, 2006) I followed Jackson and Pargetter (1973) in calling generic probabilities "indefinite probabilities", but I have never liked that terminology.

with symptoms *S*. That generic probability is entirely compatible with the singular probability of a particular person with symptoms *S* having the disease being something other than .6. For example, in addition to the symptoms, Carlos might have a family history that favors having the disease, and that might raise the probability of *his* having it to .7. The relationship between singular probabilities and related generic probabilities is the subject of the theory of *direct inference*. There are a number of different theories of direct inference (Reichenbach 1949, Kyburg 1974, Pollock 1984, 1990, 2006, 2008, Halpern 1990, Bacchus 1990, Bacchus et al 1996). They differ in details, but they all agree that the generic probability of an *S* being a *D* need not be the same as the singular probability that a particular *S* is a *D*. Applying this to the bridge and the airplane, although the generic probabilities of bridges collapsing and airplanes crashing are not 0, the singular probability of a particular bridge collapsing at a specific time or a particular airplane crashing at a certain time might be 0.

The singular probabilities that we employ in decision-theoretic reasoning are sensitive to what we believe about the framework of the decision problem, and also to general probabilistic facts about the world. What kind of probability could work like this? As far as I can see, the only way to define a probability that takes account of our beliefs about the framework is to conditionalize on those beliefs. More precisely, Pollock (1984, 1990, 2006, 2008) suggests that we start with an objective singular probability prob, which takes account of all the general facts about the world, including generic probabilities. Then relative to a cognizer *S*, we define the *mixed physical/epistemic probability* PROB(P) to be the result of conditionalizing prob on the set of *S*'s warranted propositions. That is, where **W** is the conjunction of *S*'s warranted propositions, we let PROB(P) = prob($P/\mathbf{W}$). This automatically takes account of whatever *S* is warranted in believing, and relativizes the probabilities of other propositions to *S*'s epistemic state.[12]

Why conditionalize on warranted propositions rather than either (1) justified beliefs or (2) *S*'s knowledge (Williamson 2000, Hawthorne and Stanley forthcoming). The problem with conditionalizing on the set of justified beliefs is that they might not be logically consistent. A set of beliefs could be justified for *S* even though there is a subtle inconsistency that *S* has not detected. By conditionalizing on the warranted propositions instead, we make sure that all inconsistencies have been resolved. The problem with conditionalizing on knowledge is that knowledge must be true, but *S*'s epistemic state is not determined just by his true beliefs. For example, in deciding whether to take an umbrella, when *S* looks out the window and thinks he sees rain clouds he might be fooled by a cleverly designed projector displaying an image of rain clouds on the outside of the window. That would contribute to his epistemic state, and make it reasonable to take the umbrella, but it would not contribute to his knowledge.

This is the only way I can see to make sense of a probability that attaches to propositions and is sensitive to the agent's epistemic state. If someone thinks there is another way to do this, I challenge him to produce it. But defining PROB in this way automatically assigns probability 1 to warranted propositions. Furthermore, if an agent is justified in believing *P*, he is presumably justified in believing that *P* is a warranted proposition (at least, if he thinks about it a bit), so he is justified in believing that PROB(P) = 1. I take it that this explains the bridge example and the airplane example. It also undermines the conjunctivitis objection. If we are only conjoining propositions we believe to have probability 1, we can reasonably expect arbitrarily long conjunctions of them to also have probability 1. This is not intended to be a deep point about warrant. Rather, it is a trivial consequence of what I think is the only way to construct a singular probability that takes account of the cognizer's epistemic state.

It is not clear how the mixed physical/epistemic probability PROB is related to epistemology. It automatically assigns PROB(P) = 1 for all warranted propositions *P*. For propositions whose supporting arguments are too weak to provide full-fledged justification, we might think of PROB(P) as a measure of something epistemic, although the tautology problem demonstrates that it is not degrees of justification. Perhaps it could be regarded as measuring degrees of warrant that are too weak to constitute warrant simpliciter. However, degrees of warrant are not quite what we want to be talking about when we are doing epistemology. We want to know what an agent should believe given the current state of his reasoning, not supposing (per impossible) that he has done all possible relevant reasoning.

This section has discussed several issues of interest to epistemologists, but putting it back in context, the main conclusion I want to draw is that we cannot salvage subjective probabilities for real agents by identifying them with either degrees of justification or degrees of warrant. The

[12] Pollock (2006) argues that PROB is still not quite the right probability for decision-theoretic reasoning. What we actually need is a causal probability C-PROB defined in terms of PROB.

unavoidable conclusions are (1) that there is no way to make sense of subjective probability for real resource-bounded agents, and (2) that epistemic cognition does not have a structure that can be modeled using probabilities of any kind. The elegance of the probability calculus has seduced people into embracing Bayesian epistemology, but in the end it must be forsaken for more traditional-looking epistemological theories.

9. Bayesian Conditionalization

There is an interesting residue of Bayesian epistemology that we have not yet discussed. Even if there is no way to make sense of subjective probabilities, for decision-theoretic reasoning we need some kind of probability that takes account of the agent's epistemic situation. Standard Bayesian probability theory proves unable to supply the needed epistemic probability, but I suggested that mixed physical/epistemic probability might do the trick instead. Whatever kind of epistemic probability we finally seize upon, we still have the problem of how to update these probabilities in the face of new evidence. Bayesian conditionalization was proposed as the way to update subjective probabilities, and it might be suggested that this is of more general applicability, working for whatever kind of epistemic probability is appropriate for decision-theoretic reasoning. Let us evaluate this suggestion specifically with regard to mixed physical/epistemic probability.

First, we face a problem concerning what counts as evidence. There are two possibilities. On the one hand, it might be insisted that the evidence to which Bayesian conditionalization appeals is always low-level sensory input. For vision this would consist of percepts. This is not the way practicing Bayesians normally construe evidence. However, we do not have to follow the Bayesian in this. Let us consider this construal of evidence. There are two problems with identifying evidence with percepts. First, most epistemologists agree that we do not usually attend to our percepts. When we form beliefs on the basis of perception, our beliefs are at the level of "The cat is sitting on the dinner table licking the dirty dishes," not (something crudely approximated by) "I have a visual percept containing a fuzzy multi-colored blob with a moving pink blob on one side moving rapidly in the vicinity of some white oval shapes with multi-colored blobs on their surfaces." If we do not have beliefs about our percepts, then we do not have degrees of belief about them either, and presumably do not assign probabilities to them. But then we cannot use Bayesian conditionalization to conditionalize on them, because that requires having a probability prob(E) for the evidence.

On a non-Bayesian approach to epistemic probabilities, it might be possible to assign probabilities to percepts and conditionalize on those probabilities even if we have no beliefs about them. The only way I can see to do this is to have a flat probability distribution over all possible percepts. However, percepts are psychophysical entities constructed by subdoxastic cognitive processes, and they can be described along multiple psychophysical dimensions, such as position, shape, color, texture, etc. At least some of these dimensions would seem to have infinite ranges (e.g., position and shape). This, however, requires that a flat probability distribution will assign probability 0 to all of them. And that would be fatal to Bayesian conditionalization, because by definition $\mathrm{prob}_E(P) = \dfrac{\mathrm{prob}(P \,\&\, E)}{\mathrm{prob}(E)}$, and the latter is only well-defined if $\mathrm{prob}(E) \neq 0$.

These problems are, I think, insurmountable for the proposal that the evidence employed in Bayesian conditionalization should consist of percepts. This will not bother most philosophers who are enamored of Bayesian conditionalization, because that is not what they normally take as evidence anyway. It is much more common to take evidence to consist of meter readings, or the results of sampling, or even higher-level observations like "There is an advance in the perihelion of Mercury" or "There is increased electrical activity in the amygdala". In fact, it is common to let evidence consist of any beliefs we might acquire.

Bayesian conditionalization automatically assigns probability 1 to evidence, which accords with the treatment of warranted propositions in mixed physical/epistemic probability. However, there are two kinds of situations in which Bayesian conditionalization goes wrong. First, using Bayesian conditionalization, evidence simply accumulates. For example, if we first get the evidence E_1 and then the evidence E_2, our resulting probability is $\mathsf{PROB}_{E_1\&E_2}$. It follows from the probability calculus that $\mathsf{PROB}_{E_1\&E_2}(E_1) = 1$. In other words, new evidence can never lower the probability of earlier evidence. However, if evidence need not consist of percepts, then evidence will typically be accepted on the basis of reasoning. Furthermore, the reasoning supporting substantive (logically contingent) beliefs always contains some defeasible steps (otherwise they would be supported by deductive arguments and would be logically necessary). So new evidence could support defeaters for the arguments for earlier evidence, and contrary to Bayesian conditionalization, that should lower the probability of the earlier evidence. For example, E_1 might be "x is red", justified by the fact that x looks red to the agent. But E_2 might be "x is illuminated by red lights, and red lights can make things look red when they are not". Given E_2, the probability of E_1 should no longer be 1.

A different kind of problem arises from the fact that the "conditional reliability" of non-deductive (defeasible) inference schemes is rarely 1. That is, if P is a defeasible reason for Q, it will not usually be the case that $\mathrm{PROB}(Q/P) = 1$. For example, x's looking red to S provides S with a defeasible reason for believing that x is red, but $\mathrm{PROB}(x \text{ is red}/x \text{ looks red to } S) \neq 1$. Nevertheless, if P is a defeasible reason for Q, P is warranted, and we reason from P to Q in the absence of any defeaters, that can make Q warranted. Accordingly $\mathrm{PROB}_P(Q) = 1$, although $\mathrm{PROB}(Q/P) \neq 1$. Thus $\mathrm{PROB}_P(Q) \neq \mathrm{PROB}(Q/P)$.

These examples make it evident that epistemic cognition has too rich a logical structure to be captured by simple Bayesian conditionalization.

Although Bayesian conditionalization is not a generally adequate principle for updating mixed physical/epistemic probabilities in the light of new evidence, we can formulate a principle that is reminiscent of it. Recall that, where **W** is the conjunction of S's warranted propositions, $\mathrm{PROB}(P) = \mathrm{prob}(P/\mathbf{W})$. Let E be some new evidence. As E will normally be accepted on the basis of reasoning from other propositions some of which are themselves newly accepted, we should take E to include all of this new evidence. Then let $\mathbf{W}\oplus E$ be the conjunction of propositions warranted for S when S starts with the warranted propositions **W** and then acquires the new evidence E. We can think of "$\oplus E$" as an epistemic update operator for warranted propositions. It then follows from the definition of mixed physical/epistemic probability that $\mathrm{PROB}_E(P) = \mathrm{prob}(P/\mathbf{W}\oplus E)$. This is reminiscent of Bayesian updating, however, it reduces the updating of probabilities to the epistemic updating of warrant, and it leaves the latter to a separate epistemological theory rather than assuming that $\mathbf{W}\oplus E$ is the logical closure of $\mathbf{W}\&E$.

10. Conclusions

Bayesian epistemology consists of a theory of subjective probability coupled with Bayesian conditionalization. The latter purports to tell us how subjective probabilities should be updated in the face of new evidence. Because of its simplicity and mathematical elegance, Bayesian epistemology has a seductive appeal to philosophers with a formal bent. It appears to get a great deal out of very little. However, if something seems too good to be true, it usually is. I have argued that there is no way to make sense of subjective probability for real epistemic agents. And when we substitute a more realistic kind of epistemic probability for subjective probability, it becomes clear that Bayesian conditionalization is simplistic. The general lesson to be learned is that epistemic cognition has a more complex logical structure than countenanced by Bayesian epistemology.

References

Bacchus, Fahiem
1990 Representing and Reasoning with Probabilistic Knowledge, MIT Press.
Bacchus, Fahiem, Adam J. Grove, Joseph Y. Halpern, Daphne Koller
1996 "From statistical knowledge bases to degrees of belief", *Artificial Intelligence* **87**, 75-143.
Bolker, Ethan
1966 "Functions resembling quotients of measures," *Transactions of the American Mathematical Society* **124**, 293–312.
Cherniak, Christopher
1986 *Minimal Rationality*. Cambridge, Mass, MIT Press.
de Finetti, Bruno
1974 *Theory of Probability*, vol. 1. New York: John Wiley and Sons.
Halpern, J. Y.
1990 "An analysis of first-order logics of probability", *Artificial Intelligence* **46**, 311-350.
Hawthorne, John, and Jason Stanley
forthcoming "Knowledge and action", *The Journal of Philosophy*.
Howson, Colin and Peter Urbach
1989 *Scientific Reasoning: The Bayesian Approach*. LasSalle, Ill., Open Court.

Jackson, Frank, and John Pargetter
1973 "Indefinite probability statements," *Synthese 26*, 205–215.
Jeffrey, Richard
1983 *The Logic of Decision, 2nd edition*, University of Chicago Press.
Joyce, James
1998 The Foundations of Causal Decision Theory. Cambridge, Cambridge
Kyburg, Henry, Jr.
1970 "Conjunctivitis", Marshall Swain (ed.), *Induction, Accveptance, and Rational Belief*. Dordrecht, D. Reidel.
1974 The Logical Foundations of Statistical Inference. Dordrecht: Reidel.
Lewis, David
1981 "Causal decision theory", *Australasian Journal of Philosophy* **59**, 5-30.
1994 "Chance and credence: Humean supervenience debugged," *Mind* **103**, 473–490.
Luce, R. Duncan and David Krantz
1971 "Conditional expected utility", *Econometrica* **39**, 253-71.
Plantinga, Alvin
1993 *Warrant and Proper Function.* New York: Oxford University Press.
Pollock, John
1984 The Foundations of Philosophical Semantics, Princeton University Press.
1990 Nomic Probability and the Foundations of Induction, Oxford University Press.
2006 Thinking about Acting: Logical Foundations for Rational Decision Making, New York: Oxford University Press.
2008 "Reasoning defeasibly about probabilities", in Michael O'Rourke and Joseph Cambell (eds.), *Knowledge and Skepticism*, Cambridge, MA: MIT Press.
Ramsey, Frank
1926 "Truth and probability," R. *B*. Braithwaite (ed.), *The Foundations of Mathematics*. Paterson, N.J., Littlefield, Adams.
Reichenbach, Hans
1949 *A Theory of Probability*. Berkeley: University of California Press. (Original German edition 1935)
Savage, Leonard
1954 *The Foundations of Statistics*. New York, Dover.
Skyrms, Brian
1980 *Causal Necessity*, Yale University Press, New Haven.
1984 *Pragmatics and Empiricism*, Yale University Press, New Haven.
Tversky, Amos, and David Kahneman
1974 "Judgment Under Uncertainty: Heuristics and Biases", *Science* **185**:1124-1131.
van Fraassen, Bas
1981 *The Scientific Image*. Oxford: Oxford University Press.
1985 "Empiricism in the philosophy of science", in *Images of Science: Essays on Realism and Empiricism*, eds. P. M. Churchland and C. A. Hooker, Chicago, University of Chicago Press, 245-308.
White, Roger
2006 "Problems for Dogmatism", *Philosophical Studies* **131**, 525-557.
Williamson, Timothy
2000 *Knowledge and Its Limits*, Oxford: Oxford University Press.