fallisphilosophicalstudies.pdf
4/16/02

Goldman on Probabilistic Inference*

Don Fallis
School of Information Resources
University of Arizona
fallis@email.arizona.edu

Philosophical Studies, **109**, 3, (2002): 223-240
http://dx.doi.org/10.1023/A:1019644929122
The original publication is available at www.springerlink.com.

Goldman on Probabilistic Inference

Abstract

In his latest book, *Knowledge in a Social World*, Alvin Goldman claims to have established that if a reasoner starts with accurate estimates of the reliability of new evidence and *conditionalizes* on this evidence, then this reasoner is *objectively likely* to end up closer to the truth. In this paper, I argue that Goldman's result is not nearly as philosophically significant as he would have us believe. First, accurately estimating the reliability of evidence—in the sense that Goldman requires—is not quite as easy as it might sound. Second, being objectively likely to end up closer to the truth—in the sense that Goldman establishes—is not quite as valuable as it might sound.

1. Introduction

Epistemologists are commonly interested in identifying inference procedures that have good epistemic properties. Deductive inference, for instance, has a particularly valuable epistemic property. Namely, if you start with true premises and you reason deductively, you are guaranteed to reach true conclusions. Probabilistic inference, however, by its very nature cannot provide such a guarantee.

Even so, Alvin Goldman claims to have established that a particular form of probabilistic inference has an analogous property. In particular, he claims that if you start with accurate estimates of the reliability of new evidence and you *conditionalize*[1] on this evidence, you are objectively likely to end up closer to the truth.[2] Or, as Goldman puts it in *Knowledge in a Social World*, "use of Bayes' Theorem will not always raise the user's degree of knowledge; but, under conditions to be specified, it is *objectively likely* to raise

his degree of knowledge. This is a significant property, which should certainly not be belittled" (115).[3] Indeed, if Goldman has really established what he claims to have established, it would be an extremely significant philosophical result.

Goldman's overall project in *Knowledge in a Social World* is to identify practices that have good epistemic consequences. This project is analogous to that of a moral consequentialist who wants to identify practices that have good moral consequences (87). And "Bayesian reasoning with accurate likelihoods" (122) is understandably Goldman's prime example of a practice that has good epistemic consequences. For example, he discusses its implications for the evaluation of scientific evidence (260-263), evidence presented in a court of law (292-295), and testimonial evidence in general (115-125).

In order to establish his claim about this particular form of probabilistic inference, Goldman and Moshe Shaked have proven a new mathematical theorem (121).[4] In section **2** of this paper, I explain this theorem using the framework of decision theory. In the sections that follow, however, I argue that Goldman seriously overstates the philosophical significance of this theorem. In particular, I discuss two reasons why the interest and scope of the Goldman/Shaked theorem are rather limited. First, only rarely will human beings be able to *accurately estimate*—in the precise sense required by the theorem—the reliability of evidence. Second, it is not always epistemically valuable to be objectively likely to end up *closer to the truth*—in the precise sense established by the theorem.

As I indicate below, the Goldman/Shaked theorem does have some interesting philosophical implications, but it simply will not bear the full philosophical weight that Goldman wants to put on it. However, the interest of the present paper extends beyond

the status of Goldman's claim about this particular form of probabilistic inference. An analysis of Goldman's claim illustrates some potential pitfalls of epistemic consequentialism in general that should be avoided.

2. Decision Theoretic Epistemology

Let *h* be the proposition that Floyd Thursby murdered Miles Archer (see Hammett 1934). And, suppose that Sam Spade wants to determine whether or not *h* is true. Or, to use Goldman's terminology, suppose that Spade wants to increase his "degree of knowledge" with respect to *h*. In other words, Spade wants to raise his degree of belief in *h* if *h* is true and lower his degree of belief in *h* if *h* is false (90).[5]

In an attempt to achieve this goal, Spade might try gathering new evidence about the murder. For example, he could check to see if Thursby's fingerprints are on the murder weapon.[6] However, it only makes sense to check for Thursby's fingerprints if Spade expects to increase his degree of knowledge with respect to *h* by doing so.

In order to decide what to do, Spade needs to compare the *expected epistemic payoff* (EEP) of checking for Thursby's fingerprints with the EEP of not checking.[7] In order to do this, Spade needs to know (a) the possible outcomes for each course of action, (b) the epistemic payoff for each of these outcomes, and (c) the probability of each of these outcomes.

Let *e* be the proposition that Thursby's fingerprints are on the murder weapon. There are then four possible states of the world: *h*&*e*, *h*&~*e*, ~*h*&*e*, and ~*h*&~*e*. In other words, Thursby might or might not be the murderer and his fingerprints might or might not be on the murder weapon.

Let $P_S(h)$ be Spade's initial degree of belief in Thursby's guilt. And, suppose (as I will throughout this paper) that Spade conditionalizes on new evidence. Thus, if he checks for Thursby's fingerprints and gets the result *e*, his new degree of belief in Thursby's guilt will be the probability that Spade assigns to Thursby being the murderer given that his fingerprints are on the murder weapon ($P_S(h|e)$). If he gets the result ~*e*, his new degree of belief in Thursby's guilt will be $P_S(h|\sim e)$. If he does not check for Thursby's fingerprints, his degree of belief in Thursby's guilt will continue to be $P_S(h)$.

The epistemic payoff for each possible outcome is simply Spade's degree of belief in the *truth* about *h*.[8] Thus, if *h* is true, the epistemic payoff is Spade's degree of belief in *h*. So, for example, if Spade's degree of belief in Thursby's guilt is 0.8 and Thursby really is the murderer, the epistemic payoff to Spade is 0.8. If *h* is false, the epistemic payoff is one minus Spade's degree of belief in *h* (i.e., it is Spade's degree of belief in ~*h*).

	h&*e*	*h*&~*e*	~*h*&*e*	~*h*&~*e*
check	$P_S(h\|e)$	$P_S(h\|\sim e)$	$P_S(\sim h\|e)$	$P_S(\sim h\|\sim e)$
don't check	$P_S(h)$		$P_S(\sim h)$	

Given that Spade conditionalizes on new evidence, we can use Bayes' Theorem to calculate these epistemic payoffs as long as we know (a) Spade's initial degree of belief in Thursby's guilt ($P_S(h)$), (b) the probability that Spade assigns to Thursby's fingerprints being on the murder weapon given that Thursby is the murderer ($P_S(e|h)$), and (c) the probability that Spade assigns to Thursby's fingerprints being on the murder weapon given that Thursby is not the murderer ($P_S(e|\sim h)$). In other words, these epistemic payoffs can be calculated just using Spade's initial degree of belief in Thursby's guilt and Spade's estimate of the reliability of this fingerprint evidence.

In addition, it is possible to calculate the probability that Spade assigns to each of the outcomes just using Spade's initial degree of belief in Thursby's guilt and Spade's estimate of the reliability of this fingerprint evidence.[9] And, once we have calculated the epistemic payoffs for each outcome and the probabilities of these outcomes, it is a simple matter to calculate the EEP for each course of action. If the EEP of checking for Thursby's fingerprints is greater than the EEP of not checking, then Spade should expect to increase his degree of knowledge by checking for Thursby's fingerprints.

Let us look at a specific example. Suppose that $P_S(h) = 0.8$, $P_S(e|h) = 0.75$, and $P_S(e|{\sim}h) = 0.25$.[10] Plugging these numbers into Bayes' Theorem, we find that the epistemic payoffs are as follows:

	h&*e*	*h*&~*e*	~*h*&*e*	~*h*&~*e*
check	0.92	0.57	0.08	0.43
don't check	0.80		0.20	

The probabilities that Spade assigns to the four outcomes are 0.60, 0.20, 0.05, and 0.15 respectively. Hence, the EEP of checking for Thursby's fingerprints is 0.73 (i.e., 0.60·0.92 + 0.20·0.57 + 0.05·0.08 + 0.15·0.43). This is higher than the expected epistemic payoff of not checking which is only 0.68 (i.e., 0.80·0.80 + 0.20·0.20). Thus, in this particular case, Spade should expect to increase his degree of knowledge by checking for Thursby's fingerprints.

Several years ago, Paul Horwich (1982, 127-129) proved that this result holds in general. That is, as long as Spade conditionalizes on the fingerprint evidence, Spade should expect to increase his degree of knowledge by checking for Thursby's fingerprints. It does not matter what Spade's initial degree of belief in Thursby's guilt is or what his estimate of the reliability of the fingerprint evidence is.

Of course, the Horwich theorem only establishes that checking for fingerprints is *subjectively* likely (i.e., from Spade's perspective) to increase Spade's degree of knowledge. Goldman's claim, however, is that checking for fingerprints is *objectively* likely to increase Spade's degree of knowledge.

In any particular case, we can determine if checking for fingerprints is *objectively* likely to increase Spade's degree of knowledge simply by calculating the expected epistemic payoffs using the *objective* probabilities of the different possible outcomes (rather than the probabilities that Spade happens to assign to these outcomes). However, since it is somewhat awkward to talk about "*objective* expectations," I will use the device of talking about "*God's* expectations." I will assume that God is omniscient (i.e., that God's probabilities *are* the objective probabilities).[11] And, in particular, I will assume that God knows for sure whether Thursby is the murderer and that God knows exactly how reliable the fingerprint evidence is.

Let us look at a specific example. Suppose that Thursby is the murderer (i.e., that $P_G(h) = 1$). Also, suppose that Spade's initial degree of belief in Thursby's guilt is 0.8 (i.e., that $P_S(h) = 0.8$). Finally, suppose that Spade accurately estimates the reliability of the fingerprint evidence. In particular, suppose that $P_G(e|h) = P_S(e|h) = 0.75$, and $P_G(e|\sim h) = P_S(e|\sim h) = 0.25$. Now, since $P_G(h) = 1$, there are really only two possible states of the world from God's perspective: $h\&e$, $h\&\sim e$. And, using Spade's initial degree of belief in Thursby's guilt and his estimate of the reliability of the evidence, we can calculate the epistemic payoff for each of the possible outcomes:

	h&e	*h&~e*
check	0.92	0.57
don't check	0.80	

Hence, from God's perspective, the EEP of Spade checking for Thursby's fingerprints is approximately 0.83 (i.e., 0.75·0.92 + 0.25·0.57). This is higher than the epistemic payoff of 0.80 that God expects if Spade does not check. Thus, God will expect Spade to increase his degree of knowledge by checking for Thursby's fingerprints. In fact, we would reach the same conclusion if we were to suppose instead that Thursby is *not* the murderer.

The Goldman/Shaked theorem establishes that this result holds in general. That is, as long as Spade accurately estimates the reliability of the fingerprint evidence (and Spade conditionalizes on this evidence), God will expect Spade to increase his degree of knowledge by checking for Thursby's fingerprints.[12] It does not matter what Spade's initial degree of belief in Thursby's guilt is. It also does not matter whether Thursby is the murderer. All that matters is that Spade's estimate of the reliability of the fingerprint evidence matches God's estimate of the reliability of the fingerprint evidence. Thus, Goldman concludes that, if you start with accurate estimates of the reliability of new evidence and you conditionalize on this evidence, you are *objectively* likely to end up closer to the truth.

Another interesting consequence of the Goldman/Shaked theorem is that *any* individual whose estimate of the reliability of the fingerprint evidence matches Spade's estimate will expect Spade to increase his degree of knowledge by checking for Thursby's fingerprints. It does not matter whether this individual's degree of belief in Thursby's guilt is anywhere near Spade's degree of belief in Thursby's guilt.

3. Objective Probabilities

There is, however, something strange about the example that was used to illustrate the Goldman/Shaked theorem in the previous section. It was assumed that God—being omniscient—knows for sure whether Thursby is the murderer. However, it was also assumed that God does *not* know for sure whether Thursby's fingerprints are on the murder weapon. In particular, we assumed that God only assigned a probability of 0.75 to Thursby's fingerprints being on the murder weapon. But, if God is omniscient, wouldn't God know for sure whether or not Thursby's fingerprints are on the murder weapon? In other words, wouldn't the objective probability of Thursby's fingerprints being on the murder weapon be either zero or one?

Even if the objective probabilities are all either zero or one, the Goldman/Shaked theorem is still true. That is, it is still the case that if Spade accurately estimates the reliability of the fingerprint evidence, then God will expect Spade to increase his degree of knowledge by checking for Thursby's fingerprints. However, as I explain below, if the objective probabilities are all either zero or one, the interest and scope of the theorem is rather limited.

If the objective probabilities are all either zero or one, then—in order for Spade to accurately estimate the reliability of the fingerprint evidence—he would have to know for sure whether Thursby's fingerprints would be on the murder weapon if he were the murderer. And, he would have to know for sure whether Thursby's fingerprints would be on the weapon if he were not the murderer. However, this is a degree of knowledge that human beings are unlikely to attain or even approach. In addition, if Spade is required to have this degree of knowledge, then the theorem is trivial to prove because Spade will

raise his degree of belief in Thursby's guilt—on the basis of the fingerprint evidence—if and only if Thursby is the murderer.[13]

Of course, we may have been too quick in concluding on the basis of God's omniscience that objective probabilities are all either zero or one. Maybe Einstein was wrong and God *does* play dice with the universe. And unless the universe is completely deterministic, there are at least some events that are *chancy*. For instance, the objective probability that a particular fair coin will land heads when I flip it five minutes from now is somewhere between zero and one (viz., the objective probability is 0.5). However, it turns out that this chanciness does not greatly increase the interest or scope of the Goldman/Shaked theorem.[14]

For one thing, chanciness only applies to events that are still in the future (see, e.g., Beebee and Papineau 1997, 221). For example, the objective probability that a particular fair coin landed heads when I flipped it five minutes ago is either zero or one. Either it landed heads or it did not. Similarly, by the time that Spade is ready to check for Thursby's fingerprints, they are either on the murder weapon or they are not. As a result, the probability of the fingerprints being there given that Thursby is the murderer is either zero or one.[15] In general, by the time you are ready to update your degrees of belief on the basis of a particular piece of evidence, the objective probability of that piece of evidence is one.[16]

In any case, even with respect to chancy events, human beings will rarely be able to determine the objective probabilities. For example, suppose that the murder has not yet occurred so that the objective probability of Thursby's fingerprints being on the murder weapon given that Thursby is the murderer is still somewhere between zero and

one.[17] Now, Spade may very well be able to find out that 75% of murderers leave their fingerprints on the murder weapon. However, the objective probability of *Thursby's* fingerprints being on the murder weapon given that *Thursby* is the murderer need not be anywhere near 0.75. For example, Thursby may have a habit of wearing gloves. And it may be the case that only 5% of murderers who habitually wear gloves leave their fingerprints on the murder weapon. In fact, there may be all sorts of unique properties of Thursby that bear on the objective probability of *his* fingerprints being on the murder weapon. Unless Spade can discover everything about Thursby that bears on this probability, he cannot determine what this probability is.

Of course, given that Spade does not know that Thursby has a habit of wearing gloves (or anything else about Thursby that bears on the probability of his fingerprints being on the murder weapon), there is a sense in which 0.75 is the correct probability for Spade to assign to Thursby's fingerprints being on the murder weapon given that Thursby is the murderer. In fact, we might say that 0.75 is the objective probability—*relative to Spade's state of knowledge*—of Thursby's fingerprints being on the murder weapon given that Thursby is the murderer (cf. Beebee and Papineau 1997, 223-224). However, even if Spade's probabilities are correct in this sense, he still may be objectively likely (i.e., from God's perspective) to *decrease* his degree of knowledge by checking for Thursby's fingerprints (e.g., if God is sure that Thursby's fingerprints are not on the murder weapon even though he is the murderer).

It should be noted that there are certainly some cases where it is feasible to determine the objective probability of a chancy event. For example, we can determine the objective probability that a particular fair coin will land heads when I flip it five

minutes from now. However, it does not seem to be feasible to precisely determine objective probabilities in most of the everyday cases where we are trying to increase our degree of knowledge. For example, it does not seem feasible in one of Goldman's favorite applications of the Goldman/Shaked theorem: viz., the evaluation of evidence presented in a court of law. Murderers just do not exhibit the high degree of uniformity that fair coins do.[18]

Of course, it may not actually be necessary for Spade to precisely determine the reliability of the fingerprint evidence in order to be objectively likely to end up closer to the truth. It might be sufficient, for example, for Spade's estimate of the reliability of the fingerprint evidence to simply approximate the actual reliability. Indeed, the Goldman/Shaked theorem does not establish that starting with accurate estimates of the reliability of new evidence and conditionalizing on this evidence is the *only* way to be objectively likely to increase one's degree of knowledge.[19]

Even so, Goldman has nothing definitive to tell us about any other ways to be objectively likely to increase one's degree of knowledge. Goldman does conjecture that "with a fixed prior, the prospect for improving one's degree of knowledge of the target proposition is better with close approximations to the likelihoods, or their ratios, than with poor approximations" (123). However, strangely enough, for many initial degrees of belief in Thursby's guilt, even greater increases in Spade's degree of knowledge would be objectively likely if Spade were to *overestimate* the reliability of the fingerprint evidence.[20] Thus, Goldman's conjecture is false.

The Goldman/Shaked theorem does identify certain conditions under which gathering new evidence is objectively likely to increase one's degree of knowledge. In

this section, I have argued that these specific conditions are rarely going to be satisfied. In the next section, I argue that (even if these conditions are satisfied) it is not necessarily a good idea to do things that are objectively likely to increase one's degree of knowledge.

4. Epistemic Risk

Goldman's epistemic consequentialism is a normative project. He is not simply interested in describing the epistemic consequences of various practices. He wants to conclude that certain practices are epistemically *meritorious* based on the fact that they lead to *good* epistemic consequences. In particular, he wants to conclude that "Bayesian reasoning with accurate likelihoods … is a veritistically meritorious practice" (122) based on the fact that it is objectively likely to increase one's degree of knowledge. Unfortunately, the fact that a practice is objectively likely to increase one's degree of knowledge is not necessarily a good thing. In particular, even if *X*—Spade, God, etc.—thinks that the EEP of checking for Thursby's fingerprints is greater than the EEP of not checking, *X* will not necessarily think that it is a good idea for Spade to check.

Now, as Goldman himself notes (6), there may be non-epistemic costs that outweigh the epistemic benefits of checking for fingerprints. For example, checking for fingerprints on the murder weapon might be prohibitively expensive. However, even if we focus our attention on the purely epistemic issue of closeness to the truth, *X* still might not think that it is a good idea for Spade to check for Thursby's fingerprints.

Consider an analogous case. Spade is offered a choice between $1 million for sure and a 50% chance at $4 million. In this case, the expected monetary payoff of the gamble is much greater than the expected monetary payoff of the sure thing. However,

even if all that Spade cares about is money, he still might not think that it is a good idea to take the gamble. In other words, while the expected monetary payoff of the gamble is much greater than the expected monetary payoff of the sure thing, the expected *utility* of the gamble is probably somewhat less than the expected *utility* of the sure thing. This is because, like most other people, Spade is somewhat averse to monetary risk. And, just as one can be averse to monetary risk, one can be averse to epistemic risk (see, e.g., Levi 1967, 86-90).[21] René Descartes, for example, was clearly averse to epistemic risk in his philosophical inquiries.

In order for *X* to think that it is a good idea for Spade to check for Thursby's fingerprints, *X* must think that the expected epistemic *utility* (EEU) of checking is greater than the EEU of not checking. Now, if *X*'s epistemic utility function is linear (i.e., if *X* is risk neutral with respect to Spade's closeness to the truth), then the EEU of checking is greater than the EEU of not checking whenever the EEP of checking is greater than the EEP of not checking. However, if *X*'s epistemic utility function is concave (i.e., if *X* is averse to Spade taking epistemic risks), then the EEU of checking will often be less than the EEU of not checking even when the EEP of checking is greater than the EEP of not checking.

Let us look at a specific example. Suppose that $P_S(h) = 0.8$, $P_S(e|h) = 0.75$, and $P_S(e|\sim h) = 0.25$. Also, suppose that $P_G(h) = 1$, $P_G(e|h) = 0.75$, and $P_G(e|\sim h) = 0.25$. In section **2** of this paper, using these same numbers, we determined that God thinks the EEP of checking is greater than the EEP of not checking. However, suppose that God is somewhat averse to Spade taking epistemic risks.[22] In particular, suppose that the

epistemic utility of a degree of belief of x in the truth is equal to $2x - x^2$ rather than x. That is, suppose that God has the following concave epistemic utility function:

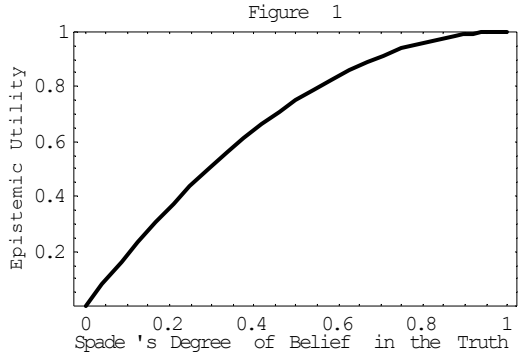

Using this concave epistemic utility function, the epistemic utilities of the various possible outcomes are as follows:

	h&*e*	*h*&~*e*
check	0.99	0.82
don't check	0.96	

Hence, from God's perspective, the EEU of Spade checking for Thursby's fingerprints is only 0.95 (i.e., 0.75·0.99 + 0.25·0.82). This is less than the epistemic utility of 0.96 that God expects if Spade does not check for Thursby's fingerprints. Thus, even though he does think that the EEP of checking is greater than the EEP of not checking, God will think that it is *not* a good idea for Spade to check for Thursby's fingerprints.

It should be noted that bringing attitudes toward epistemic risk into the discussion does not shift our attention beyond the issue of closeness to the truth. Attitudes toward epistemic risk simply determine how the epistemic benefit of moving closer to the truth compares with the epistemic benefit of not falling further into error. And, as long as the objective probabilities of the various possible states of the world are somewhere between

zero and one, these attitudes toward epistemic risk can have an effect on whether Spade should check for Thursby's fingerprints.

The Goldman/Shaked theorem does identify certain conditions under which gathering new evidence is objectively likely to increase one's degree of knowledge (as defined in section **2**). And, if one happens to be perfectly risk neutral with respect closeness to the truth, then one will consider this to be a good epistemic consequence. However, if one has some other attitude toward epistemic risk, one will not necessarily consider this to be a good epistemic consequence. And, there are many different possible attitudes toward epistemic risk.[23] In the absence of compelling reasons to think that epistemic utility functions must be linear, the Goldman/Shaked theorem does not establish that "Bayesian reasoning with accurate likelihoods … is a veritistically meritorious practice." No such reasons are given in *Knowledge in a Social World*.[24]

5. The Problem with Risk Neutrality

Goldman has not made the case that people should be risk neutral with respect to closeness to the truth. But even worse, there is good reason to think that people should *not* be perfectly risk neutral with respect to closeness to the truth. Namely, if we adopt a linear epistemic utility function, we will think that it is a good idea to engage in a practice that is clearly not epistemically meritorious. In particular, while people with linear epistemic utility functions will think that it is a good idea to gather and conditionalize on new evidence, they will also think that it is a good idea to simply jump to conclusions.[25]

For example, suppose that Spade's epistemic utility function is linear. Also, suppose once again that Spade's current degree of belief in Thursby's guilt is 0.8. In that

case, the epistemic utility of Spade's current belief state is 0.8 if Thursby is guilty. And, the epistemic utility of Spade's current belief state is 0.2 if Thursby is innocent. Now, since Spade thinks that there is an 80% chance that Thursby is guilty and a 20% chance that Thursby is innocent, he currently thinks that there is an 80% chance that the epistemic utility of his current belief state is 0.8 and a 20% chance that the epistemic utility of his current belief state is 0.2. Thus, as we discovered in section **2**, from Spade's perspective, the EEU of his current belief state is 0.68 (i.e., 0.8·0.8 + 0.2·0.2).

	h	~*h*
stay with a degree of belief of 0.8	0.8	0.2
switch to a degree of belief of 1	1	0

However, if Spade were to be *sure* that Thursby was guilty (i.e., if his degree of belief in Thursby's guilt were 1), then the epistemic utility of his belief state would be 1 if Thursby is guilty and it would be 0 if Thursby is innocent. Now, Spade currently thinks that there is an 80% chance that Thursby is guilty and a 20% chance that Thursby is innocent. Thus, from Spade's perspective, the EEU of being sure that Thursby is guilty is 0.8 (i.e., 0.8·1 + 0.2·0) which is higher than the EEU of his current belief state. As a result, Spade can increase his EEU simply by jumping to the conclusion that Thursby is guilty. In general, people with linear epistemic utility functions can always increase their EEU simply by assigning a probability of one to the hypothesis that they currently think is most likely to be true.

Admittedly, this only establishes that jumping to this conclusion is subjectively likely (i.e., from Spade's perspective) to increase Spade's EEU. It is only objectively likely to increase Spade's EEU if he is, in fact, going to jump the right way. However, this fact should not provide much comfort to the epistemic consequentialist. Spade, of

course, can only base his decisions on the probabilities that he currently assigns to the events in the world.

Now, Goldman might try to defend the use of linear epistemic utility functions by appealing to *doxastic involuntarism* (see, e.g., Alston 1988). That is, he might claim that it is not psychologically possible to simply change one's degree of belief in the absence of any new evidence. In other words, jumping to a conclusion in this way is just not an option for Spade. However, this sort of move should not provide any more comfort to the epistemic consequentialist than it would to the moral consequentialist.

For example, suppose that a particular moral consequentialist claims that *reducing the amount of suffering in the world* is our moral objective. She then concludes that various practices are morally praiseworthy because they tend to further this moral objective. We might, however, attempt to criticize her position on the grounds that it has unacceptable implications. In particular, suppose that there is a lone individual living in an isolated cave in Antarctica who is in constant pain from the unforgiving cold. As a result, there is another practice (viz., painlessly killing this lone individual) that would also tend to reduce the amount of suffering in the world, but that is clearly morally objectionable. It would not be an effective defense of her position for the moral consequentialist to merely point out that, given his remote location, it is just not possible for anyone to actually kill this individual.[26]

Finally, it should be noted that this problem with linear epistemic utility functions is not a general problem for epistemic consequentialism. There are epistemic utility functions that do not imply that jumping to conclusions is a good idea. For example, the concave epistemic utility function discussed in the previous section is such a function.[27]

Unfortunately, as we saw in the previous section, if we adopt this particular epistemic utility function, there are some circumstances under which it is not a good idea—from God's perspective—for Spade to check for Thursby's fingerprints. In other words, even if you start with accurate estimates of the reliability of new evidence and you conditionalize on this evidence, gathering this evidence is sometimes objectively likely to lead to a decrease in epistemic utility.[28]

6. Conclusion

Goldman and Shaked have proved an interesting mathematical theorem. And, strictly speaking, this theorem does show that if you start with accurate estimates of the reliability of new evidence and you conditionalize on this evidence, you are objectively likely to end up closer to the truth. However, only on rare occasions will individuals be able to accurately estimate—in the sense required by the theorem—the reliability of evidence. Also, being objectively likely to end up closer to the truth—in the sense established by the theorem—is only going to be a good epistemic consequence if we happen to be risk neutral with respect to closeness to the truth. And, in addition, there is good reason to think that we should not to be risk neutral with respect to closeness to the truth. As a result, while the claim that "Bayesian reasoning with accurate likelihoods … is a veritistically meritorious practice" is almost certainly correct, it is not clear how much support the Goldman/Shaked theorem provides for this claim.[29]

References

Albert, David. 1992. *Quantum Mechanics and Experience*. Cambridge: Harvard.

Alston, William P. 1988. "The Deontological Conception of Epistemic Justification." *Philosophical Perspectives* 2:257-99.

Beebee, Helen and David Papineau. 1997. "Probability as a Guide to Life." *Journal of Philosophy* 94:217-43.

Conee, E. and R. Feldman. 1998. "The Generality Problem for Reliabilism." *Philosophical Studies* 89:1-29.

Fallis, Don and Gerrard Liddell. "Further Results on Inquiry and Truth Possession." unpublished manuscript.

Goldman, Alvin. 1999. *Knowledge in a Social World*. New York: Oxford.

Goldman, Alvin and Moshe Shaked. 1991. "An Economic Model of Scientific Activity and Truth Acquisition." *Philosophical Studies* 63:31-55.

Good, I. J. 1974. "A Little Learning Can Be Dangerous." *British Journal for the Philosophy of Science* 25:340-342.

Hammett, Dashiell. 1934. *The Maltese Falcon*. New York: Modern Library.

Horwich, Paul. 1982. *Probability and Evidence*. Cambridge: Cambridge.

Levi, Isaac. 1967. *Gambling With Truth*. Cambridge: MIT.

Maher, Patrick. 1990. "Why Scientists Gather Evidence." *British Journal for the Philosophy of Science* 41:103-19.

Oddie, Graham. 1986. *Likeness to Truth*. Dordrecht: Reidel.

Rawls, John. 1971. *A Theory of Justice*. Cambridge: Harvard.

Notes

[*] I would like to thank Brad Armendt, Andrew Cohen, Alvin Goldman, John Lewis, Peter Lewis, Kay Mathiesen, Glenn Ross, Joseph Shieber, Brian Skyrms, and an anonymous referee for helpful comments and suggestions.

[1] That is, the probability that you assign to a proposition after observing the evidence is equal to the conditional probability that you assigned—prior to observing the evidence—to the proposition given the evidence.

[2] As I explain in section **2**, the phrase "you will end up closer to the truth" here just means that you will have a higher degree of belief in a true proposition. It does not mean that the proposition that you believe will have greater *verisimilitude* (see, e.g., Oddie 1986).

[3] Unless otherwise noted, references are to Goldman 1999.

[4] This theorem first appeared in Goldman and Shaked 1991.

[5] An individual's degree of belief in the true answer to a "target question" (117) is what Goldman calls the "veritistic value" (88) of that individual's current belief state (with respect to that question). There are epistemic values other than closeness to the truth—simplicity, explanatory power, justification, etc. However, in this paper, I follow Goldman (23-25) and simply treat knowledge as true belief.

[6] In *Knowledge in a Social World*, Goldman introduces the Goldman/Shaked theorem as part of a discussion of *testimonial* evidence (115-125). However, the theorem is intended to apply generally to any type of evidence (260-263).

[7] Using decision theory to analyze *epistemic* decisions is not a new idea (see, e.g., Levi 1967).

[8] In this paper, I will take the epistemic payoff to be Spade's degree of knowledge with respect to *h* rather than the *change* in his degree of knowledge (120). The results come out the same either way. In addition, following Goldman, I will calculate the EEP with respect to a single proposition of interest, *h*. Of course, finding out if Thursby's fingerprints are on the murder weapon will typically affect Spade's degree of belief in many other propositions (e.g., the proposition that Thursby habitually wears gloves). As a result, the EEP of checking for Thursby's fingerprints would ideally be calculated with respect to Spade's total belief state.

[9] For example, $P_S(h\&e) = P_S(h) \cdot P_S(e|h)$.

[10] $P_S(e|h)$ and $P_S(e|\sim h)$ sum to 1 in this particular example, but they need not.

[11] Evaluating expected payoffs from the perspective of an omniscient being is not a new idea (see, e.g., Good 1974).

[12] Of course, Spade might actually *lower* his degree of knowledge by checking for Thursby's fingerprints (if, e.g., $h\&\sim e$ happens to be the actual state of the world).

[13] In other words, he is not just objectively likely to increase his degree of knowledge in this case; he is *certain* to increase his degree of knowledge.

[14] The points that I discuss below apply whether the chanciness is due to physical indeterminacy (as in the case of a coin flip) or due to the free will of human beings.

[15] If Thursby is not the murderer, then this conditional probability is undefined.

[16] Under some interpretations of quantum mechanics, some events in the past remain chancy until a measurement is made (see, e.g., Albert 1992). However, by the time you are ready to conditionalize on an event, the relevant measurement has presumably been made.

[17] It should be noted that we are interested in the objective probability of this particular piece of evidence being left by this particular murderer. As Goldman explains with regard to eyewitness testimony, we are interested in "the chance that the witness would be inaccurate on *this* particular occasion, concerning the specific proposition being asserted" (124).

[18] The relative frequency of fingerprints being left by a particular class of murderers may be fairly stable. However, taking the objective probability of Thursby's fingerprints being on the murder weapon given that Thursby is the murderer to be a relative frequency is not a promising strategy. The problem here is analogous to the generality problem for process reliabilism (see Conee and Feldman 1998). As I have noted above, Thursby—if he is the murderer—belongs to several different classes of murderers. Thus, unless one of these classes is privileged in some way, this strategy does not necessarily pin down a definite objective probability.

[19] Goldman sometimes seems to suggest otherwise. For example, he states that "if a probabilistic reasoner begins with inaccurate likelihoods, she cannot expect the Bayesian method to improve her V-value vis-à-vis the target question" (117).

[20] For example, suppose that $P_S(h)$ is somewhere between 0.26 and 0.74. Also, suppose that $P_G(e|h) = 0.75$ and $P_G(e|\sim h) = 0.25$. In that case, regardless of whether Thursby is the murderer, the expected increase in Spade's degree of knowledge is greater if $P_S(e|h) = 0.8$ and $P_S(e|\sim h) = 0.2$ (i.e., if he overestimates the reliability of the evidence) than if $P_S(e|h) = 0.75$ and $P_S(e|\sim h) = 0.25$ (i.e., if he accurately estimates the reliability of the evidence).

[21] In addition, an individual's attitude toward epistemic risk might vary for different propositions.

[22] It should be noted that we are assuming once again—for the sake of argument—that God does not know for sure whether Thursby's fingerprints are on the murder weapon.

[23] Basically, any monotonically increasing function is a potential epistemic utility function.

[24] Somewhat like Goldman, John Rawls wanted to claim—in *A Theory of Justice*—that a particular attitude toward risk should be adopted in a particular context. (Specifically, he wanted to claim that the individuals in the original position should be extremely risk averse.) However, unlike Goldman, Rawls offered some support for the claim that this was the correct attitude toward risk in this particular context (see, e.g., Rawls 1971, 172).

[25] This problem with linear epistemic utility functions was first pointed out by Patrick Maher (1990, 112-113).

[26] We will suppose this individual is not able kill himself for psychological reasons.

[27] The epistemic utility functions proposed by Levi (1967) are also such functions.

[28] In fact, this is true—with only one exception—for any epistemic utility function that does not imply that jumping to conclusions is a good idea (see Fallis and Liddell). The one exception is when epistemic utility goes to $-\P$ as degree of belief in the truth goes to 0 (e.g., when the epistemic utility of a degree of belief of x in the truth is equal to $\log(x)$). Thus, in order to take advantage of this exception, you have to be willing to take the rather extreme position that—*epistemically speaking*—it is *infinitely* bad for Spade to be completely wrong about Thursby's guilt.

[29] It should be noted that Miles Archer was, in fact, murdered by Brigid O'Shaughnessy (see Hammett 1934).