

fallisepistemtwo.pdf
6/7/06

Epistemic Value Theory and Social Epistemology

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

forthcoming in [*Episteme*](#), 2, 3, (2006): 177-188

Epistemic Value Theory and Social Epistemology

Abstract: In order to guide the decisions of real people who want to bring about good epistemic outcomes for themselves and others, we need to understand our *epistemic values*. In *Knowledge in a Social World*, Alvin Goldman has proposed an *epistemic value theory* that allows us to say whether one outcome is epistemically better than another. However, it has been suggested that Goldman's theory is not really an *epistemic value theory* at all because whether one outcome is epistemically better than another partly depends on our *non-epistemic* interests. In this paper, I argue that an epistemic value theory *that serves the purposes of social epistemology* must incorporate non-epistemic interests in much the way that Goldman's theory does. In fact, I argue that Goldman's theory does not go far enough in this direction. In particular, the epistemic value of having a particular true belief should actually be *weighted* by how interested we are in the topic.

Men are likely to forget in the heat of action where their best interest lies.
– Caspar Gutman, Esq.

Introduction

People typically seek information (e.g., by attending schools, visiting libraries, and accessing the Internet) in order to bring about good *epistemic* outcomes. That is, they want to acquire knowledge, true beliefs, justified beliefs, understanding, etc. *Epistemology* is primarily the study of what such outcomes amount to (e.g., what does it mean to say that *S* is justified in believing *p*?). However, there has recently been a lot of discussion of exactly why the sorts of outcomes studied by epistemologists are valuable and of exactly how valuable they are (cf. Kvanvig 2003, Riggs 2006).

This interest in *epistemic values* has been motivated largely by perennial questions in epistemology, such as why knowledge is more valuable than true belief (cf. Plato 1961, *Meno*, 98a). *Social epistemology*, however, provides an independent motivation for studying epistemic values.¹ Namely, we need to understand epistemic values in order to explain or to guide the decisions of real people who want to bring about good epistemic outcomes for themselves and others.

What an agent decides to do depends largely on what that agent considers to be a good outcome (cf. Resnik 1987). Thus, when we attempt to explain or to guide the decisions that people make, we have to determine what their *values* are. For instance, many philosophers of science (e.g., Levi 1962, Maher 1993, Fallis 1997) have tried to determine what the epistemic values (or *epistemic utilities*) of scientists must be in order to explain their methodological choices.

¹ More generally, *applied* epistemology (whether social or individual) provides this same motivation (cf. Bishop and Trout 2005).

Alvin Goldman (1999) recently has developed a theory to guide the decisions of real people who seek good epistemic outcomes in many other social arenas. According to Goldman, the project of social epistemology is to evaluate practices in the law, the media, education, government, science, etc. in terms of their epistemic consequences. But in order to determine whether a practice has good epistemic consequences (or whether one practice has better epistemic consequences than another), we have to be able to say when one outcome is epistemically better than another. In other words, we need an *epistemic value theory*.²

The epistemic value theory that Goldman has developed in order to evaluate social practices has been criticized on a number of points (cf. Levy 2000, Maffie 2000, Schmitt 2000, Feldman 2001, Fallis 2002, Kitcher 2002). Most notably, his theory has been criticized on the grounds that it does not incorporate many important epistemic values, such as having *understanding*, having a *coherent* belief system, having beliefs that are in line with the *evidence*, and having *reasons* for one's beliefs (cf. Levy 2000, 283, Feldman 2001, 166). But more intriguingly, it has also been suggested that his theory is not really an *epistemic* value theory at all because it *does* incorporate *non-epistemic* values in a central way (cf. Maffie 2000, 251-252, Feldman 2001, 165-166).³ In this paper, I will argue that an epistemic value theory *that serves the purposes of social epistemology* must incorporate non-epistemic values in much the way that Goldman's theory does.

Goldman's epistemic value theory

The two epistemic values most often discussed by epistemologists are *having true beliefs* and *avoiding error* (cf. James 1979, 24, Alston 1985, 59, Latus 2000). Goldman's (1999, 87-100) epistemic value theory incorporates only these two epistemic values. Basically, true beliefs make an outcome epistemically better and false beliefs make an outcome epistemically worse.⁴ It is for this reason that Goldman actually refers to his theory as a *veritistic* value theory.

Goldman's theory might be extended to incorporate other epistemic values. But Goldman (1999, 100) is intentionally offering a "relatively simple and uncluttered framework" in the hope that such a framework "will smooth the path to analytical progress." Since the worry about non-epistemic values would still arise even if his theory did incorporate other epistemic values, these two main epistemic values will be the exclusive focus of this paper. In fact, following Goldman (1999, 24-25), I will use the term "knowledge" in this paper in the weak sense of true belief.⁵

² Paterson (1979) also developed an epistemic value theory to guide curriculum decisions in educational contexts.

³ No one has yet been able to draw a bright line to distinguish epistemic values from non-epistemic values (cf. Riggs 2006). But some values (e.g., having true beliefs) are clearly epistemic and some values (e.g., avoiding pain) are clearly non-epistemic. This is all that I will need for purposes of this paper.

⁴ In this paper, I will focus on *categorical beliefs*, but epistemic value theory can also be extended to cover *degrees of belief* (cf. Goldman 1999, 90, Fallis 2002).

⁵ This is purely a matter of stylistic convenience. I am not here endorsing Crispin Sartwell's (1992) view that knowledge merely is true belief.

In constructing an epistemic value theory, we restrict our focus to one specific type of value that people have (cf. Riggs 2006). We leave other types of value (e.g., ethical values, financial values, nutritional values) to other specialists (cf. Goldman 1999, 6). Having true beliefs and avoiding error are clearly epistemic values. But non-epistemic values also turn out to play a significant role in Goldman's epistemic value theory. Goldman (1999, 89) claims that the epistemic value of having a true belief (and avoiding error) should only "be assessed relative to *questions of interest*."⁶ For example, there is no epistemic value to having a true belief about a subject (e.g., the exact number of grains of sand on the local beach) that one has no interest in.

Strictly speaking, the interests that play this role in his epistemic value theory can be epistemic as well as non-epistemic. For example, particular questions may be of interest because finding out correct answers to them may lead to the acquisition of more true beliefs in the future (cf. Feldman 2001 166-167).⁷ However, the specific issue that Goldman has been criticized for is having *non-epistemic* interests play such a significant role in his epistemic value theory. And that is the issue that I will focus on here.

In this paper, I will argue that non-epistemic interests should play a significant role in epistemic value theory. When we try to develop an epistemic value theory to help real people make decisions with epistemic consequences, it is just not possible to completely bracket the non-epistemic interests that people have. As Goldman (1999, 95) puts it, "a good epistemic practice should not be indifferent to the relative amounts of interest that a given agent takes in different questions."

In fact, I also argue that Goldman does not go far enough in this direction. Goldman (1999, 95) wants non-epistemic interests to play only a "moderate" rather than a "pervasive" role in his epistemic value theory. However, I argue that non-epistemic interests should actually play a more pervasive role in his theory. Namely, the epistemic value of having a particular true belief should be *weighted* by how interested one is in the topic (cf. Kitcher 2002, 192-193).

On having true beliefs

Goldman (1999, 89-90) tends to focus on cases where we are interested in having a correct answer to a single question. For example, in a criminal trial, our ultimate epistemic goal is that the jury acquire a true belief about the guilt of the defendant. We may have other epistemic goals in this context, but they are always a means to the end of acquiring this single true belief.

⁶ Goldman (1999, 94-95) actually discusses three types of interest that might play a role in epistemic evaluation. An individual might be actively interested in the answer to a question, she might be interested if she thought about the question, or she might be interested if she were better informed. Frederick Schmitt (2000, 270-273) suggests some additional types of interest that might be incorporated into Goldman's epistemic value theory. But these distinctions will not matter for the purposes of this paper.

⁷ Finding out *incorrect* answers to some questions may also lead to the acquisition of more true beliefs in the future (cf. Nozick 1993, 69).

However, we typically want correct answers to many questions. This is certainly the case in the domain of social epistemology. For example, libraries and schools are places where we want many people to be able to acquire many true beliefs on many topics (cf. Paterson 1979, Goldman 1999, 349-373, Fallis 2006). In such contexts, we have to determine how to allocate our resources across topics (e.g., when we make acquisition decisions in libraries or curriculum decisions in schools). In other words, we have to determine whether having one set of true beliefs is epistemically better than having another set of true beliefs. Here is a simple example:

There are many true beliefs that would be of great interest to us, such as a true belief (b_1) about “how the world could be made a good and happy place” (Adams 1981). However, there are also many true beliefs that would be of absolutely no interest to us. Nozick’s (1993, 67) example is a true belief (b_2) about “the precise number of grains of sand on each beach in the world.” One of Goldman’s (1999, 88) examples is a true belief (b_3) about “the 323rd entry in the Wichita, Kansas, telephone directory.” With all of this in mind, which of the following outcomes is epistemically better?

Outcome **A** (i.e., believing b_1):

number of true beliefs
1

Outcome **B** (i.e., believing b_2 and b_3):

number of true beliefs
2

Clearly, we would much rather know just b_1 than know both b_2 and b_3 . Libraries, for example, are ready to expend more resources helping people to acquire knowledge that is of interest to them than they are ready to expend helping people to acquire knowledge that is not. However, if we just focus on the *epistemic* value of having true beliefs, it seems that we might have to say that outcome B is epistemically better than outcome A because it involves more true beliefs.

It should be noted that the foregoing discussion makes the idealizing assumption that beliefs can be individuated and counted. However, there are clearly serious difficulties with individuating and counting beliefs (cf. Latus 2000, 30-31). For example, consider a belief about who the 323rd person in the telephone directory is and a belief about what the last name of the 323rd person in the telephone directory is. Do these really count as two separate beliefs? But, in the interest of “analytical progress,” I am simply going to assume (for purposes of this paper) that beliefs can be individuated and counted.⁸

Now, it might be suggested that we should really focus on the *information content* of the true beliefs in an outcome rather than on the *number* of true beliefs (cf. Maher 1993, 213, Schmitt 2000, 272). This is certainly a legitimate move, but it is not clear that it will allow us to conclude that outcome A is epistemically better than outcome B. That is, the

⁸ As noted above, Goldman tends to focus on the epistemic value of having single beliefs. As a result, he is largely able to avoid the issue of individuating and counting beliefs.

information content of the true belief in outcome A is not necessarily greater than the information content of the true beliefs in outcome B. In fact, if we are lucky, a true belief about “how the world could be made a good and happy place” could turn out to be quite simple.⁹ In any event, we could easily modify this example by adding true beliefs about other trivial matters to outcome B until their cumulative information content exceeded the information content of b_1 . But we would presumably still rather know b_1 than know even an extremely large number of trivial facts.

It might also be suggested that outcome A could be epistemically better than outcome B because it does better with respect to epistemic values other than truth. However, the description of these two outcomes could easily be filled out so that outcome A is equivalent to outcome B with respect to these other epistemic values. For example, we could have just as good reasons for believing b_2 and b_3 as we have for believing b_1 . But even if they were equivalent with respect to these other epistemic values, we would presumably still prefer outcome A to outcome B.

Fortuitously, Goldman’s theory does provide us with a way to reach the desired conclusion about this case. According to Goldman (1999, 89), true beliefs do not all have the same epistemic value. In particular, having a true belief has epistemic value only if it is an answer to a question of interest (cf. Maher 1993, 214). Thus, outcome A and outcome B should really be represented as follows (with true beliefs multiplied by 1 if they are of interest and by 0 if they are not):

Outcome A:

value of true belief
$1 \mu 1 = 1$

Outcome B:

value of true belief
$2 \mu 0 = 0$

Thus, according to Goldman’s theory, outcome A is epistemically better than outcome B given that we are not interested in knowing b_2 or b_3 .

Now, on Goldman’s theory, *non-epistemic* interests can be an important part of what determines which outcome is *epistemically* better. In particular, the reason that outcome A is epistemically better than outcome B is that it is simply much more *useful* to know b_1 than it is to know both b_2 and b_3 . But, as I will now argue, this is an epistemically acceptable role for non-epistemic interests to play in an epistemic value theory. Non-epistemic interests simply serve to pick out one ordering, from a set of *epistemically*

⁹ In the case that he describes, Douglas Adams (1981) seems to suggest that the true belief in question is fairly simple: “One day, nearly two thousand years after one man had been nailed to a tree for saying how great it would be to be nice to people for a change, a girl sitting on her own in a small café in Rickmansworth suddenly realized what it was that had been going wrong all this time and she finally knew how the world could be made a good and happy place. This time it was right, it would work, and no one would have to get nailed to anything.”

acceptable orderings, as the one that should guide the decisions of a person who seeks good epistemic outcomes.

There is an important sense in which “outcome A is epistemically better than outcome B” and “outcome B is epistemically better than outcome A” are both epistemically acceptable orderings. Both orderings satisfy the following normative constraint:

NC1: All other things being equal (epistemically speaking), a particular person *not* having a particular true belief is never epistemically better than that person having that true belief.¹⁰

In other words, being ignorant (or in error) about a particular topic is never epistemically better than having a true belief about it. This is a fairly weak constraint on orderings. For example, it does not tell us whether having knowledge about topic X is epistemically better than having knowledge about topic Y. In fact, it does not even tell us whether having *more* knowledge about topic X is epistemically better than having *less* knowledge about topic Y.¹¹ But it is nevertheless a substantive constraint. There are many orderings of outcomes that fail to satisfy it. For example, it would be epistemically unacceptable to say that knowing just b_2 is epistemically better than knowing both b_2 and b_3 .

Basically, the normative constraints of epistemology only determine a *partial* ordering of outcomes in terms of epistemic betterness. In other words, there are several different complete (i.e., *linear*) orderings that are epistemically acceptable.¹² But when making decisions, it is often important that we have a complete ordering of outcomes in mind (cf. Resnik 1987, 23). For example, a student may have to decide between pursuing knowledge about topic X and pursuing knowledge about topic Y. Picking out such a complete ordering is essentially the role that non-epistemic interests play in Goldman’s epistemic value theory.

Now, it is certainly possible to imagine many ways in which non-epistemic interests could play an *epistemically unacceptable* role in an epistemic value theory. For example, our epistemic value theory might say that non-epistemic interests can *trump* the value of having a true belief. In fact, our theory might even say that non-epistemic interests can determine *which* beliefs are true (cf. James 1975, 97). But non-epistemic interests do not play such roles in Goldman’s theory.¹³ For instance, suppose that having a true belief

¹⁰ The *ceteris paribus* clause is necessary because it might be epistemically better for this person not to have a particular true belief *and* not to have a bunch of false beliefs, for example. But the *ceteris paribus* clause has to be restricted to the epistemic realm because it would not be *epistemically* better for this person not to have a particular true belief *and* not to be in terrible pain, for example.

¹¹ If epistemology did tell us that having more knowledge about one topic is always epistemically better than have less knowledge about another topic, then we would have to accept that knowing an extremely large number of trivial facts is always epistemically better than just knowing b_1 .

¹² In this paper, I will just appeal to NC1 and NC2 (see below). There may very well be even tighter normative constraints on having true beliefs (and on avoiding error). However, as Patrick Maher (1993, 210) suggests, it seems unlikely that such constraints will determine a complete ordering of outcomes.

¹³ Non-epistemic interests do not *trump* our epistemic values on Goldman’s theory, but we might say that they do *shape* our epistemic values.

(b₄) that her son has committed a horrible crime would make a mother's life miserable (cf. Nozick 1993, 69). In that case, it probably is better for her not to have this true belief, or even for her to believe the opposite. But this fact does not make it true that her son is innocent nor does it make it *epistemically* better for her to believe that he is.

Finally, it is important to emphasize that Goldman's proposal only shows that outcome A is epistemically better than outcome B *if* the person making a decision is not interested in trivial facts about grains of sand and telephone books. If this person were interested in such facts and not interested in how to make the world a better place, then outcome B would be epistemically better than outcome A. In order to conclude that outcome A is epistemically better than outcome B under all circumstances, it would be necessary to develop what might be called a *thick* theory of epistemic value.

Several epistemologists (e.g., Paterson 1979, Kitcher 2004, 216, Bishop and Trout 2005, 93-103) have tried to develop theories where the epistemic value of having a particular true belief depends on its *objective importance* (rather than on anyone's interest in the topic).¹⁴ R. W. K. Paterson (1979, 95), for example, thinks that different pieces of knowledge should be assigned different values based on their "cognitive richness" and "objective worth." And it is certainly plausible that knowledge about "how the world could be made a good and happy place" is objectively more important than knowledge about grains of sand or telephone books (cf. Nozick 1993, 67-68). Thus, a theory along these lines would presumably determine an ordering of outcomes such that outcome A is epistemically better than outcome B under all circumstances.

The development of such a thick theory of epistemic value is certainly an interesting project for epistemologists to pursue. But it is not clear that a thick theory would serve to guide the decisions of the vast majority of people who seek the truth (and try to avoid error). The decisions of such people can be epistemically acceptable even if they do not seek the truth about the objectively most important matters. For example, it is plausible to think that facts about physics have greater "cognitive richness" and "objective worth" than facts about sociology. But this does not mean that it is epistemically unacceptable to major in sociology. Even if she is going to acquire less knowledge on a less important topic, a sociology major still prefers having true beliefs to being ignorant or in error.¹⁵ In fact, even a person who is deeply interested in trivial facts about grains of sand and telephone books prefers having true beliefs to being ignorant or in error. While it may be unwise for an individual to have such interests, it is not clear that it is *epistemically* unacceptable.¹⁶ Thus, it is still important to have a *thin* theory of epistemic value (of the sort that Goldman has proposed) to guide the decisions of real people.¹⁷

On distributing true beliefs among people

¹⁴ Ernest Sosa (2003, 158-159) questions whether it is possible to develop such a theory.

¹⁵ Science as a whole would certainly be in trouble if nobody studied physics. But this would not mean that the sociology majors did not seek good epistemic outcomes.

¹⁶ Compare Hume's (2000, 267) claim that it is "not contrary to reason to prefer the destruction of the whole world to the scratching of my finger."

¹⁷ In a similar vein, Maher (1993, 209-216) has essentially proposed a thin theory of the epistemic values of *scientists*.

Just as true beliefs can be distributed over different topics, they can be distributed among different people.¹⁸ For example, the various crew members on a ship will typically know different things about sailing the ship (cf. Hutchins 1995).¹⁹ Some crew members will know about navigation, other crew members will know about maintenance, etc. And non-epistemic interests play the same sort of role here in determining which distributions of true beliefs are epistemically better than others. And the distribution that involves the greatest number of true beliefs is again not necessarily the epistemically best outcome.

In general, there will be several orderings of distributions that satisfy the aforementioned normative constraint on orderings (NC1). Non-epistemic interests simply pick out one of these epistemically acceptable orderings of distributions. For example, the epistemically best distribution of knowledge among the various crew members is probably one that allows the ship to function most efficiently (cf. Fallis 2004, 107-109).²⁰ As Goldman (1999, 96) puts it, “a veritistically good practice for such an enterprise would promote the required distribution of knowledge, even if that does not translate into a high average knowledge across the whole team.”

On having somewhat interesting true beliefs

Goldman’s theory allowed us to reach the desired conclusion that outcome A is epistemically better than outcome B. There are many other cases, however, where Goldman’s theory does not allow us to reach the desired conclusion. Here is a simple example:

There are many true beliefs that would be of some, but not very great, interest to us. For example, it is of interest to us to have a true belief (b_5) about what time *Hill Street Blues* comes on TV or a true belief (b_6) about what time *Streets of San Francisco* comes on TV. Even so, we would much rather know just b_1 than know both b_5 and b_6 . However, on Goldman’s theory, it seems that we have to say that outcome C is epistemically better than outcome A because it involves more true beliefs that are of interest.

Outcome C (i.e., believing b_5 and b_6):

value of true belief
$2 \mu 1 = 2$

There is an obvious extension to Goldman’s theory, however, that does allow us to reach the desired conclusion. We simply *weight* true beliefs by one’s *degree* of interest (cf.

¹⁸ Like other goods, knowledge can be distributed over times as well as among people (cf. Broome 1991, 22-25). The same sorts of points made here about how knowledge should be distributed among people can also be made about how knowledge should be distributed over times (cf. Fallis 2004, 111-113).

¹⁹ The requisite knowledge for sailing a ship will often be knowledge of *how* to do something rather than knowledge *that* something is the case (cf. Snowdon 2003). But Goldman’s epistemic value theory only explicitly addresses propositional knowledge.

²⁰ It might be the case that the ship runs most efficiently when certain crew members do not know certain things (e.g., for security reasons). In that case, the best distribution *all things considered* would not be the epistemically best distribution.

Goldman 1999, 95, Kitcher 2002, 192). For example, suppose that we are one million times more interested in knowing b_1 than in knowing either b_5 or b_6 . In that case, outcome A should really be represented as follows:

Outcome A:

value of true belief
$1 \mu 1,000,000 = 1,000,000$

Thus, this weighting proposal allows us to reach the desired conclusion that outcome A is epistemically better than outcome C.²¹

Goldman (1999, 95), however, explicitly rejects this weighting proposal.²² He is worried that “this would lead down the path toward abandoning the specialized, veritistic mission of epistemology in favor of a more purely pragmatic enterprise.” However, it is not clear that this weighting proposal really does go too far down this path. For example, the weighting proposal does not say that non-epistemic interests can give *negative* weight to true beliefs. Like Goldman’s original theory, the weighting proposal satisfies the aforementioned normative constraint (NC1) that simply having fewer true beliefs never makes things epistemically better. Of course, this proposal does say that non-epistemic interests can give different weight to different true beliefs. However, Goldman himself has already taken the biggest step in this direction by saying that some true beliefs have *zero* epistemic value because of non-epistemic interests.

On avoiding error

In addition to the epistemic value of having true beliefs, there is also the epistemic value of avoiding error. As William James (1979, 24) put it, “we must know the truth; and we must avoid error—these are our first and great commandments as would-be knowers; but they are not two ways of stating an identical commandment, they are two separable laws.” Goldman’s theory can be used (in the obvious way) to determine the epistemic *disvalue* of having false beliefs. Also, his theory is consistent with a second normative constraint on orderings:

²¹ The full story of epistemic value will actually be somewhat more complicated because the epistemic value of having true beliefs is not necessarily *additive* (cf. Oddie 2001, 318). For example, the epistemic value of knowing b_5 and b_6 is not necessarily the sum of the epistemic value of knowing just b_5 and the epistemic value of knowing just b_6 . As a result, each combination of true beliefs (and false beliefs) probably needs to be weighted independently. But I will ignore this complication for purposes of this paper.

²² While he explicitly rejects the weighting proposal, Goldman (1999, 96) does say that “sensitivity to relative amounts of interest should play a modest role in assessing a practice’s epistemic credentials.” However, he does not provide any details about how such an alternative to the weighting proposal is supposed to work (cf. Kitcher 2002, 192-193). More recently, Goldman (2002, 218) has proposed an “extended veritistic epistemology” that does allow true beliefs to be weighted by degree of interest. However, it also allows non-epistemic interests to give negative weight to having true beliefs (as in the case of the mother with the criminal son). In other words, it allows for values that conflict with NC1. Such an “extended veritistic epistemology” is potentially useful because people may often have such values. However, such an epistemic value theory clearly abandons the “veritistic mission of epistemology.”

NC2: All other things being equal (epistemically speaking), a particular person having a particular false belief is never epistemically better than that person not having that false belief.

However, the sort of argument given in the previous section indicates that it is necessary to use the weighting proposal in this case as well.

On having true beliefs and avoiding error

Of course, many practices have an effect both on how many true beliefs and on how many false beliefs people acquire. Thus, in order to determine whether such a practice has good epistemic consequences, we have to be able to compare epistemic outcomes that involve both true beliefs and false beliefs. This is another area in which Goldman’s value theory does not sufficiently incorporate non-epistemic interests.

Outcome **D**:

number of true beliefs	number of false beliefs
9	2

Outcome **E**:

number of true beliefs	number of false beliefs
5	3

Suppose that all of the beliefs in outcome D and outcome E are answers to questions of interest. Also, suppose that all of the true beliefs have the same epistemic value and that all of the false beliefs have the same epistemic disvalue. Under these assumptions, it is clear that outcome D is epistemically better than outcome E since it includes more true beliefs and fewer false beliefs.

When there are two or more values, however, it is not always this easy to determine which outcome is better. It is often the case that one outcome does better with respect to one value, but worse with respect to another value. In the domain of social epistemology, this sort of conflict can easily arise between the epistemic values of having true beliefs and avoiding error. Consider, for example, the decision about whether to retain the old edition of an encyclopedia in a library collection (cf. Fallis 2006, 501-502). Retaining this edition may lead to more false beliefs because it is out-of-date, but doing so may also lead to more true beliefs (e.g., about what people believed at the time that this edition was published). Here is a simple example of this sort of conflict:

Outcome **F**:

number of true beliefs	number of false beliefs
7	1

While it is clear that outcome D is epistemically better than outcome E, it is much harder to say whether outcome D is epistemically better than outcome F. Outcome D does better with respect to having true beliefs, but outcome F does better with respect to avoiding

error. In order to resolve this sort of conflict, we need a way to aggregate the epistemic value of the true beliefs and the epistemic disvalue of the false beliefs in any given outcome into a single number that measures the *overall epistemic value* of that outcome.

Goldman's value theory provides a way to determine the overall epistemic value of an outcome.²³ He (1999, 89) says that the epistemic value of having a true belief is the same magnitude as the epistemic disvalue of having a false belief (cf. Fallis 2002). Thus, in order to determine the overall epistemic value of an outcome, we simply subtract the number of false beliefs from the number of true beliefs. For example, the overall epistemic value of outcome D is $(9 \mu 1) - (2 \mu 1) = 7$ and the overall epistemic value of outcome F is $(7 \mu 1) - (1 \mu 1) = 6$.²⁴ Thus, using Goldman's method of aggregation, outcome D is epistemically better than outcome F.

Unfortunately, Goldman gives no argument for why the value of having true beliefs is the same as the value of avoiding error. In fact, it seems perfectly reasonable to assign much greater value to avoiding error (cf. Hume 1977, 111, Descartes 1996, 12).²⁵ For example, it might be suggested that the disvalue of having a false belief is three times the value of having a true belief. In that case, the overall epistemic value of outcome D is $(9 \mu 1) - (2 \mu 3) = 3$ and the overall epistemic value of outcome F is $(7 \mu 1) - (1 \mu 3) = 4$. Thus, using this method of aggregation, outcome F is epistemically better than outcome D.

There is an important sense in which "outcome D is epistemically better than outcome F" and "outcome F is epistemically better than outcome D" are both epistemically acceptable orderings. Both orderings satisfy the aforementioned normative constraints on orderings (NC1 and NC2). Again, the normative constraints of epistemology only determine a partial ordering of outcomes (cf. Levi 1962, 57, Fallis 2004, 111). Assigning equal value to having true beliefs and avoiding error (as Goldman does) yields one epistemically acceptable complete ordering. But many other epistemically acceptable complete orderings result from assigning other relative values to having true beliefs and avoiding error

As Keith DeRose's (1992) well-known bank example indicates, different people in different contexts can have different epistemic values (e.g., based on the cost of being

²³ As noted above, Goldman tends to focus on the epistemic value of having single beliefs. In the case of a single belief, this aggregation problem does not arise. However, since Goldman (1999, 92) also wants to evaluate practices in terms of the *average* epistemic value that they lead to, he is committed to the possibility of such aggregation.

²⁴ This calculation assumes that E-value(true belief) = 1, E-value(withholding judgment) = 0, and E-value(false belief) = -1. Goldman actually says that V-value(true belief) = 1, V-value(withholding judgment) = 0.5, and V-value(false belief) = 0. However, my definition is equivalent to Goldman's definition because value functions are closed under linear transformations (cf. Broome 1991, 74). Note that E-value(x) = 2 μ (V-value(x) - 0.5). For these simple examples, it is just more convenient to take withholding judgment (rather than false belief) as the zero point.

²⁵ In fact, Riggs (2003, 347) claims that almost all epistemologists think that the value of avoiding error is greater than the value of having true beliefs. But at least a few epistemologists (e.g., James 1979, 31-32, Godfrey-Smith 1991) do think that there are some circumstances where the value of having true beliefs is greater.

wrong).²⁶ In particular, different people may legitimately assign different relative values to having true beliefs and avoiding error. Also, the same people may legitimately assign different relative values under different circumstances. For example, a person might very well assign greater value to avoiding errors when she is seeking medical information than when she is seeking sports information. Thus, non-epistemic interests can help to determine which complete ordering of outcomes should guide the decisions of a person who seeks good epistemic outcomes.

Fortuitously, the weighting proposal described above shows how non-epistemic interests can pick out one of the epistemically acceptable complete orderings. In order to determine the value of the true beliefs in an outcome, we simply weight the true beliefs by degree of interest. We use the same sort of procedure to determine the disvalue of the false beliefs. Finally, in order to determine the overall epistemic value of an outcome, we simply subtract the disvalue of the false beliefs from the value of the true beliefs.²⁷

On changing people's interests

The main objection to incorporating non-epistemic interests into an epistemic value theory is that we seem to abandon the “veritistic mission of epistemology” by doing so. In this paper, I have argued that this is not the case. But Richard Feldman (2001) has also suggested that, as a result of incorporating non-epistemic interests, Goldman's theory has some very counter-intuitive implications. And it appears to have these implications whether or not we extend Goldman's original theory with the weighting proposal discussed above.

Feldman (2001, 166) asks us to:

consider a practice that gives lots of people accurate information about what's on television each night. This enables a lot of people to know some things that they do care about. This practice has a lot of veritistic value. We can add to its value by making even more people interested in this topic and decreasing the amount of thought they give to complex matters that they are apt to get wrong. Thus, a practice that encourages such interests and distributes such information widely turns out to be epistemically better than one that encourages careful thought about social policies or complex intellectual matters.

As Feldman suggests, it would be very strange for an epistemic value theory to imply that getting people interested in trivial matters has better epistemic consequences than helping people acquire knowledge about more significant matters.

²⁶ The bank example supports this conclusion whether or not one rejects *contextualism* about knowledge.

²⁷ There may very well be some constraints on possible weightings beyond what is required by NC1 and NC2. For example, it is hard to see how an outcome that involves 9 true beliefs and 2,000,000 false beliefs could be epistemically better than outcome F. But, on the weighting proposal, it could turn out to be epistemically better if the two additional true beliefs were sufficiently interesting.

It might even be suggested that there is a more counter-intuitive implication of Goldman's epistemic value theory along these same lines. Namely, a practice that simply gets people more interested in things that they already know (and does not lead to any additional knowledge) might seem to have better epistemic consequences than a practice that leads people to know more things. However, Goldman's epistemic value theory does not have this implication. Goldman's theory is consistent with the aforementioned normative constraint (NC1) that simply having fewer true beliefs never makes things epistemically better. And since the first practice here leads to strictly fewer true beliefs, saying that the first practice has better epistemic consequences than the second practice would conflict with this constraint. Of course, even if Goldman's theory does not have this more counter-intuitive implication, it still seems strange that simply changing people's interests would make an outcome epistemically better at all.

In any event, the two practices in Feldman's actual example both lead to additional knowledge. It is just that the additional knowledge provided by the first practice is trivial while the additional knowledge provided by the second practice is much more significant. Thus, saying that the first practice in Feldman's example has better epistemic consequences than the second practice would not conflict with the aforementioned normative constraint (NC1). Even so, it is still possible to argue that Goldman's epistemic value theory does not have the counter-intuitive implication that Feldman attributes to it.

In order to defend Goldman's theory, we might argue that it is impossible, or at least prohibitively expensive, to change people's interests in this way. It is certainly possible to get people interested in different things. Advertisers do it all the time. But it might be rather difficult to get people deeply interested in topics that they currently believe (with good reason) to be trivial. Even so, this contingent fact would not provide a very effective defense of Goldman's theory.²⁸ Changing people's interests in this way would be epistemically questionable even if it were a trivial matter to do so.

However, a more effective defense can be mounted by noticing that Feldman's objection to Goldman's theory is essentially the same as an objection that has been raised against preference-satisfaction (or "want-satisfaction") versions of utilitarianism (cf. Barry 1989). Preference-satisfaction utilitarianism basically says that it is good for people to get what they want. But it can often be fairly expensive to provide people with the things that they want. It can be much more cost-effective to get them to want things that are cheaper to provide (or even to be satisfied with things that they already have). But it seems problematic if a moral theory recommends that we simply change people's preferences rather than satisfy the preferences that they actually have (cf. Sen 1992, 55).

In response to this objection, Brian Barry (1989, 281) argues that it is just a mistake to think that preference-satisfaction utilitarianism has anything to say about what preferences people should have. Preference-satisfaction utilitarianism says that it is good

²⁸ In a similar vein, if it is a consequence of an epistemic value theory that we should jump to conclusions without evidence, the mere fact that we are psychologically incapable of doing so does not provide a very effective defense of the theory (cf. Fallis 2002, 236).

for people to get what they want. But it does not say that it is good for people to want what they get.²⁹ Thus, it is silent on the matter of whether it is good to change people's preferences in various ways.

We can give the very same response to Feldman's objection. Namely, it is a mistake to think that Goldman's epistemic value theory has anything to say about what interests people should have. His theory says that it is epistemically good for people to have a lot of true beliefs about matters that are of interest to them. But it does not say that it is epistemically good for people to be interested in matters about which they have (or can easily acquire) a lot of true beliefs. Thus, it is silent on the matter of whether it is epistemically good to change people's interests in various ways.

Since it is silent on this matter, Goldman's epistemic value theory does not have the counter-intuitive implication that Feldman attributes to it. But this is not to say that an epistemic value theory should be silent on the issue of how people's interests ought to change. Determining what topics people should be interested in is arguably an important project for epistemologists to pursue. It is essentially the project of developing a thick theory of epistemic value (as described above). But the fact that Goldman does not provide such a theory does not show that there is anything wrong with the thin theory of epistemic value that he has provided.

Conclusion

If we want to make decisions that have good epistemic outcomes for ourselves and others, we have to be able to say whether one outcome is epistemically better than another. In other words, we need an epistemic value theory. Interestingly, on the theory that Goldman recently has proposed, whether one outcome is epistemically better than another can partly depend on our *non-epistemic* interests. As strange as this might seem to an epistemologist, it is actually necessary—and epistemically acceptable—for an epistemic value theory that serves the purposes of social epistemology to incorporate non-epistemic interests in this way. In fact, Goldman's theory does not really go far enough in this direction. The epistemic value of having a particular true belief should actually be *weighted* by how interested we are in the topic.³⁰

References

Adams, Douglas. 1981. *The Hitchhiker's Guide to the Galaxy*. London: BBC.

Alston, William P. 1985. "Concepts of Epistemic Justification." *Monist* 68:57-89.

Barry, Brian. 1989. "Utilitarianism and Preference Change." *Utilitas* 1:278-82.

²⁹ Admittedly, the noted philosopher Sheryl Crow *has* endorsed the view that "it's not having what you want, it's wanting what you've got."

³⁰ I would like to thank Alvin Goldman, Kay Mathiesen, Fred Schmitt, Dennis Whitcomb, and an audience at the 2005 meeting of the *New Jersey Regional Philosophical Association*, for many helpful suggestions.

Bishop, Michael A. and J. D. Trout. 2005. *Epistemology and the Psychology of Human Judgment*. Oxford: Oxford.

Broome, John. 1991. *Weighing Goods*. Cambridge, Massachusetts: Basil Blackwell.

Descartes, René. 1996. *Meditations on First Philosophy*. ed. John Cottingham. Cambridge: Cambridge.

DeRose, Keith. 1992. "Contextualism and Knowledge Attribution." *Philosophy and Phenomenological Research* 52:913-29.

Fallis, Don. 1997. "The Epistemic Status of Probabilistic Proof." *Journal of Philosophy* 94:165-86.

Fallis, Don. 2002. "Goldman on Probabilistic Inference." *Philosophical Studies* 109:223-40.

Fallis, Don. 2004. "Epistemic Value Theory and Information Ethics." *Minds and Machines* 14:101-17.

Fallis, Don. 2006. "Social Epistemology and Information Science." Pp. 475-519 in *Annual Review of Information Science and Technology*, vol. 40, ed. Blaise Cronin. Medford, New Jersey: Information Today.

Feldman, Richard. 2001. Review of *Knowledge in a Social World* by Alvin Goldman. *British Journal for the Philosophy of Science* 52:163-68.

Godfrey-Smith, Peter. 1991. "Signal, Decision, Action." *Journal of Philosophy* 88:709-22.

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

Goldman, Alvin. 2002. "Reply to Commentators." *Philosophy and Phenomenological Research* 64:215-27.

Hume, David. 1977. *An Enquiry Concerning Human Understanding*. ed. Eric Steinberg. Indianapolis: Hackett.

Hume, David. 2000. *A Treatise of Human Nature*. eds. David F. Norton and Mary J. Norton. Oxford: Oxford.

Hutchins, Edwin. 1995. *Cognition in the Wild*. Cambridge: MIT.

James, William. 1975. *Pragmatism*. Cambridge: Harvard.

- James, William. 1979. *The Will to Believe and Other Essays in Popular Philosophy*. Cambridge: Harvard.
- Kitcher, Philip. 2002. "Veritistic Value and the Project of Social Epistemology." *Philosophy and Phenomenological Research* 64:191-98.
- Kitcher, Philip. 2004. "The Ends of the Sciences." Pp. 208-29 in *The Future for Philosophy*, ed. Brian Leiter. Oxford: Oxford.
- Kvanvig, Jonathan L. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge.
- Latus, Andrew. 2000. "Our Epistemic Goal." *Australasian Journal of Philosophy* 78:28-29.
- Levi, Isaac. 1962. "On the Seriousness of Mistakes." *Philosophy of Science* 29:47-65.
- Levy, Robert. 2000. Review of *Knowledge in a Social World* by Alvin Goldman. *Teaching Philosophy* 23:279-83.
- Maffie, James. 2000. "Alternative Epistemologies and the Value of Truth." *Social Epistemology* 14:247-57.
- Maher, Patrick. 1993. *Betting on Theories*. New York: Cambridge.
- Nozick, Robert. 1993. *The Nature of Rationality*. Princeton: Princeton.
- Oddie, Graham. 2001. "Axiological Atomism." *Australasian Journal of Philosophy* 79:313-32.
- Paterson, R. W. K. 1979. "Towards an Axiology of Knowledge." *Journal of Philosophy of Education* 13:91-100.
- Plato. 1961. *Collected Dialogues*. eds. Edith Hamilton and Huntington Cairns. Princeton: Princeton University Press.
- Resnik, Michael D. 1987. *Choices*. Minneapolis: University of Minnesota Press.
- Riggs, Wayne D. 2003. "Balancing Our Epistemic Ends." *Nous* 37:342-52.
- Riggs, Wayne D. 2006. "The Value Turn in Epistemology." in *New Waves in Epistemology*, eds. Vincent F. Hendricks and Duncan Pritchard. Aldershot: Ashgate.
- Sartwell, Crispin. 1992. "Why Knowledge Is Merely True Belief." *Journal of Philosophy* 89:167-80.

Schmitt, Frederick F. 2000. "Veritistic Value." *Social Epistemology* 14:259-80.

Sen, Amartya. 1992. *Inequality Reexamined*. Cambridge: Harvard.

Snowdon, Paul. 2003. "Knowing How and Knowing That: A Distinction Reconsidered." *Proceedings of the Aristotelian Society* 104:1-29.

Sosa, Ernest. 2003. "The Place of Truth in Epistemology." Pp. 155-79 in *Intellectual Virtue*, eds. Michael DePaul and Linda Zagzebski. Oxford: Oxford.