

Epistemic Value

The topic of epistemic value has always been a central topic within epistemological discussion. Part of the reason for focussing on this topic has been a question, which can be traced right back to Plato's *Meno*, regarding what (if anything) it is about knowledge that makes it more valuable than mere true belief. Interest in the topic of epistemic value has seen a re-emergence in recent years, partly in response to a rediscovery of the *Meno* problem regarding the value of knowledge (e.g., Kvanvig 2003) and partly in response to a concern that contemporary accounts of knowledge are unable to explain the (putative) distinctive value of knowledge (e.g., Williamson 2000). Recent discussions of epistemic value have, however, extended beyond this specific issue to examine other related questions, such as whether distinct epistemic standings, like understanding, have any special value.

- **1. The *Meno* Problem**
- **2. Two Other Value Problems for Knowledge**
- **3. Reliabilism and the *Meno* Problem**
- **4. Virtue Epistemology and the Value Problem**
- **5. Understanding and Epistemic Value**
- **6. Other Accounts of the Value of Knowledge**
- **7. The Value of True Belief**
- **Bibliography**
- **Other Internet Resources**
- **Related Entries**

1. The *Meno* Problem

Much of the debate regarding epistemic value has focussed on the value of knowledge. This is unsurprising, since the epistemological focus, both in the contemporary literature and historically, has almost exclusively been on this notion. If knowledge is not of special value,

however, then this focus is somewhat mysterious. We will call the general question of why knowledge is valuable the *value problem*.

The question why knowledge is distinctively valuable has an important historical precedent in Plato's *Meno* in which Socrates raises the question of why knowledge is more valuable than mere true belief. Initially, we might appeal to the fact that knowledge appears to be of more practical use than true belief in order to mark this difference in value, but, as Socrates notes, this claim is far from obvious on closer inspection. After all, a true belief about the correct way to Larissa is surely of just as much practical use as knowledge of the way to Larissa—both will get us to our destination. Given that we clearly do value knowledge more than mere true belief, the fact that there is no obvious explanation of why this should be so creates a problem. We will call the issue of why knowledge is more valuable than mere true belief, the *Meno problem*.

Plato's own solution to this problem was to argue that the property distinctively possessed by knowledge is that of being 'tied-down' to the truth, like the mythical tethered statues of Daedalus which were so life-like that they were tied to the ground to ensure that they did not run away. In contrast, mere true belief, argues Plato, is apt to run away and be lost. Put more prosaically, the point being made here is that knowledge, unlike mere true belief, gives one a confidence that is not easily lost, and it is this property that accounts for the distinctive value of knowledge over mere true belief.

For example, if one knows the way to Larissa, rather than merely truly believes that such-and-such is the correct way to go, then one is less likely to be perturbed by the fact that the road, initially at least, seems to be going in the wrong direction. Mere true belief at this point may be lost, since one might lose all confidence that this is the right way to go. In contrast, if one knows that this is the right way to go, then one will be more sanguine in the light of this development, and thus will in all likelihood press on regardless (and thereby have one's confidence rewarded by getting where one needs to go).

Like most commentators, then, Plato responds to the *Meno* problem by trying to find a way to meet it head-on—i.e., by trying to find a way to show that knowledge is of more value than

mere true belief after all. He thus aims for a *non-revisionary* response to the problem, and we shall consider others below. Alternatively, of course, one could argue that the way to deal with this problem is simply to reject the intuition in play and argue that knowledge isn't of more value than mere true belief after all. Ideally, one would supplement such an account with an explanation of why knowledge might seem to be more valuable than mere true belief even though in fact it isn't. This would be a *revisionary* response to the problem. While few have found revisionary responses to the *Meno* problem attractive, analogous revisionary responses to other comparable problems have been relatively common, as we will see below.

2. Two Other Value Problems for Knowledge

While much of the focus of the discussion of the value of knowledge has tended to cluster around the *Meno* problem, there are in fact two further related problems in this regard. The first is what we might call—following Duncan Pritchard (2007a: §2)—the *secondary value problem* (with the *Meno* problem as the *primary value problem* for knowledge). Whereas the *Meno* problem concerns the question of why knowledge is more valuable than mere true belief, the secondary value problem concerns the issue of why knowledge is more valuable than any proper subset of its parts. That is, why do we specifically desire knowledge rather than any epistemic standing that falls short of knowing (including, but not restricted to, mere true belief)? The importance of this distinction between the two value problems can be brought out by considering a possible response to the primary value problem which is not thereby a response to the secondary value problem.

Suppose, for example, that knowledge is justified true belief plus some additional component that deals with Gettier-style cases. Suppose further, however, that justification adds value to a mere true belief. If this last point is right, then one might reasonably argue that the fact that knowledge entails justification offers a way of dealing with the primary value problem, since there would now be a property of knowledge which mere true belief lacks and which affords greater value to knowledge over mere true belief. It would not follow, however, that we would thereby have a response to the secondary value problem. This is because justified true belief is a proper subset of knowledge on our present suppositions, and thus the greater value of knowledge over mere true belief would not translate into a greater value of knowledge

over any proper subset of its parts, including justified true belief.

If one's account of the value of knowledge ended at this point, one would thus be offering a non-revisionary response to the *Meno* problem while simultaneously offering a revisionary response to the secondary value problem. Indeed, this is, in effect, the line taken by Mark Kaplan (1985), who argues that the moral of the post-Gettier literature is that what is really of epistemic value is justified true belief, and not knowledge (knowledge being justified true belief plus an additional component to rule-out Gettier-style cases). Kaplan's point is that it is of no practical consequence to us whether we have Gettier-proof justified true belief (i.e., knowledge) rather than just justified true belief, and hence there is no specific reason to value knowledge over justified true belief. Moreover, Kaplan can explain why we might ordinarily have the intuition that knowledge is of special epistemic value by noting that knowledge could very easily be confused with mere justified true belief.

It seems then that if one wishes to account for the distinctive value of knowledge, one must resolve both the *Meno* and the secondary value problem. Indeed, there may even be a third value problem for knowledge in play here. After all, one could respond to the secondary value problem by arguing that knowledge is more valuable as a matter of degree than that which falls short of knowledge. It is unclear, however, whether this way of thinking about the value of knowledge can do justice to the idea that knowledge is distinctively valuable. That is, the picture that one is left with is one on which knowledge simply marks a point on a continuum of epistemic value, but on this picture it is far from clear why the focus of epistemological theorizing has been *this* point on the continuum rather than some other point (a point just before the one that knowledge marks perhaps, or one just after). Thus, one might argue that what is required is an account of why knowledge is more valuable than that which falls short of knowledge not merely as a matter of degree but of kind, what Pritchard (2007a: §2) calls the *tertiary value problem*. In effect, the challenge posed by the tertiary value problem is to explain what special kind of value enters the picture once one gets to the point on the continuum that knowledge marks.

One further point is in order before we continue. It ought to be clear that what we are seeking when we look for a response to one of these three value problems is not an account of why

knowledge is *always* more valuable than the corresponding epistemic commodity (e.g., mere true belief). To take the *Meno* problem as an illustration on this score, no-one would surely want to hold that knowledge is always of more (overall) value than mere true belief, since there are bound to be cases in which it would better for you, all things considered, to merely truly believe p than to know p (as when knowing p would kill you, say, but merely truly believing p would win you the lottery instead). However, while it is clear that the requirement laid down on successful resolutions of the various value problems for knowledge is weaker than the demand that knowledge is always more valuable than the corresponding epistemic commodity, it isn't at all clear how best to understand this weaker demand. (Obviously, what goes here for the value problems regarding knowledge applies just as equally when it comes to analogous problems that face other epistemic standings).

3. Reliabilism and the *Meno* Problem

The first contemporary wave of work on the value problem largely concerned whether this problem raised a distinctive difficulty for reliabilist accounts of knowledge—i.e., those views which essentially define knowledge in terms of true belief that arises out of reliable belief-forming processes. In particular, the claim was that reliabilism was unable even to offer an answer to the primary value problem.

A fairly clear statement of what is at issue here is given in a number of places by Linda Zagzebski (e.g., 2003; cf. DePaul 1988; 1993; Zagzebski 1996; Jones 1997; Swinburne 1999; 2000; Riggs 2002; Kvanvig 2003; Sosa 2007: ch. 4). To begin with, Zagzebski argues that the reliability of the process by which something is produced does not automatically add value to that thing, and thus that it cannot be assumed that the reliability of the process by which a true belief is produced will add value to that true belief. In defence of this claim, she offers the analogy of a cup of coffee. She claims that a good cup of coffee which is produced by a reliable coffee machine (i.e., one that regularly produces good cups of coffee) is of no more value than an equally good cup of coffee that is produced by an unreliable coffee machine.

Furthermore, Zagzebski claims that true belief is in the relevant respects like coffee: a true

belief formed via a reliable belief-forming process is no more valuable than a true belief formed via an unreliable belief-forming process. In both cases, the value of the reliability of the process accrues in virtue of its tendency to produce a certain valuable effect (good coffee/true belief), but this means that where the effect has been produced—where one has a good cup of coffee or a true belief—then the value of the product is no greater for having been produced in a reliable way.

Elsewhere in the literature, this problem has been called the “swamping problem”, on account of how the value of true belief ‘swamps’ the value of the true belief being produced in a reliable (i.e., truth-conducive) way. So expressed, the moral of the problem seems to be that where reliabilists go awry is by treating the value of the process as being solely captured by the reliability of the process—i.e., its tendency to produce the desired effect. Since the value of the effect swamps the value of the reliability of the process by which the effect was achieved, this means that reliabilism has no resources available to it to explain why knowledge is more valuable than true belief.

It’s actually not obvious that this is a problem that is specific to reliabilism. That is, it seems that if this is a *bona fide* problem then it will affect any account of the value of knowledge which has the same relevant features as reliabilism—i.e., which regards the greater value of knowledge over true belief as instrumental value, where the instrumental value in question is relative to the valuable good of true belief. Presumably, there could be non-reliabilist views that had these features.

Even granting the main elements of the swamping argument, there are moves that the reliabilist can make in response (see, e.g., Goldman & Olsson *forthcoming*). For example, it is surely open to the reliabilist to argue that the greater instrumental value of reliable true belief over mere true belief does not need to be understood purely in terms of instrumental value relative to the good of true belief. There could, for instance, be all sorts of *practical* benefits of having a reliable true belief which generate instrumental value. (Indeed, it is worth noting that line of response to the *Meno* problem sketched by Plato which we noted above seems to specifically appeal to the greater practical instrumental value of knowledge over mere true belief).

Moreover, there is reason to think that that this objection will only at best have an impact on the most extreme forms of reliabilism—i.e., those views which *simply* say that knowledge is true belief derived from a reliable process. In contrast, more refined versions of reliabilism, such as the sort of agent reliabilist account offered by John Greco (e.g., 1999; 2000), might be thought to be untouched by this sort of argument. This is because, according to agent reliabilism, it is not any sort of reliable process which is knowledge-conducive, but only those processes that are stable features of what Greco calls the agent’s “cognitive character”. The main motivation for this restriction on reliable processes is that it excludes certain kinds of reliable processes—what Greco calls “strange and fleeting processes”—which notoriously cause problems for the view (such as processes where the reliability is due to some quirk in the subject’s environment, rather than because of any cognitive trait possessed by the agent herself). Plausibly, however, one might argue that the reliable traits that make up an agent’s cognitive character have some value independently of the instrumental value they possess in virtue of being reliable (i.e., that they have some final or intrinsic value). If this is right, then this opens up the possibility that agent-reliabilists can evade the value problem that Zagzebski identifies for pure reliabilists. (Note, however, that Greco himself does not respond to the value problem in this way. His proposal will be considered in this regard in due course).

Zagzebski’s diagnosis of what is motivating this problem for reliabilism seems, however, explicitly to exclude such a counter-response. She argues that what gives rise to this difficulty is the fact that the reliabilist has signed up to a “machine-product model of belief” (see especially, Zagzebski 2003), where the product is external to the cause. It is not clear what exactly Zagzebski means by this point, but she thinks it shows that even where the reliable process is independently valuable (i.e., independently of its being reliable), it still doesn’t follow that the value of the cause will transfer to add value to the effect. Zagzebski again offers the ‘coffee’ analogy to illustrate this: even if a reliable coffee machine were independently valuable, it would not thereby confer additional value on a good cup of coffee.

Perhaps the best way to understand what Zagzebski has in mind here is to consider what she thinks *is* required in order to resolve this problem. She argues that what is needed is an

‘internal’ connection between product and cause, such as the kind of internal connection that exists between an act and its motive which is highlighted by how we explicitly evaluate actions in terms of the motives that led to them. On this picture, then, we are not to understand knowledge as a state consisting of a known belief, but rather as a state which consists of both the true belief *and* the source from which that true belief was acquired. In short, then, the problem with the machine-product model of belief that Zagzebski claims to identify is that it leads us to evaluate the state of the knowledge independently of the means by which the knowledge was acquired. If, in contrast, we have a conception of knowledge that incorporates into the very state of knowledge the way that the knowledge was acquired, then, Zagzebski argues, we can avoid this problem.

Zagzebski’s contention is that once one effects this transition away from the machine-product model of belief, one can allow that the independent value of the reliable process can ensure that knowledge, by being produced in this way, is more valuable than mere true belief. In particular, if the process by which one gained the true belief is an epistemic virtue—a character trait which Zagzebski thinks is both reliable and intrinsically valuable—then this can ensure that the value of the knowing state in this case is more valuable than any corresponding state which simply consisted of a true belief.

Other commentators in the virtue epistemology camp, broadly conceived, have put forward similar suggestions. For example, Wayne Riggs (2002) and Greco (e.g., 2002) have argued for a ‘credit’ version of virtue epistemology, according to which the agent, in virtue of bringing about the positively valuable outcome of a true belief, is due credit as a result. Rather than treating the extra value of knowledge over true belief as deriving simply from the agent’s attainment of the target true belief, however, Riggs and Greco instead argue that we should regard the agent’s knowing as the state the agent is in when she is responsible for her true belief. Only in so doing, they claim, can we answer the value problem.

Interestingly, however, other virtue epistemologists, most notably Ernest Sosa (2003), have also advocated a ‘credit’ view of this sort, yet seem to stay within the machine-product picture of belief that Zagzebski thinks is so problematic. That is, rather than analyse the state of knowing as consisting of both the true belief and its source, they regard the state of

knowing as distinct from the process, yet treat the fact that the process is intrinsically valuable as conferring additional value on any true belief so produced. With Sosa's view in mind, it is interesting to ask just why we need to analyse knowledge in the way that Zagzebski and others suggest in order to get around the value problem.

The most direct way to approach this question is by considering whether it is really true that a valuable cause cannot confer value on its effect where cause and effect are kept separate in the way that Zagzebski claims is problematic in the case of knowledge. One commentator who has objected to Zagzebski's argument by querying this claim on her part is Berit Brogaard (2007; cf. Percival 2003; Pritchard 2007*a*: §2; 2007*b*), who claims that a valuable cause can indeed confer value on its effect in the relevant cases. Brogaard claims that virtue epistemologists like Zagzebski and Riggs endorse this claim because they adhere to what she call a "Moorean" conception of value, on which if two things have the same intrinsic properties, then they are equally valuable. Accordingly, if true belief and knowledge have the same intrinsic properties (which is what would be the case on the view of knowledge that they reject), it follows that they must have the same value. Hence, it is crucial to understand knowledge as having distinct intrinsic properties from true belief before one can hope to resolve the value problem.

If one holds that there is only intrinsic and instrumental value, then this conception of value is compelling, since objects with the same intrinsic properties trivially have the same amount of intrinsic value, and they also plausibly have the same amount of instrumental value as well (at least in the same sort of environment). As Brogaard points out, however, the Moorean conception of value is problematic because—as Wlodek Rabinowicz and Toni Roennow-Rasmussen (1999; 2003) have pointed out—there seem to be objects which we value for their own sake but whose value derives from their being extrinsically related to something else that we value. That is, such objects are *finally* (i.e., non-instrumentally) valuable without thereby being intrinsically valuable.

The standard example in this regard is Princess Diana's dress. This would be regarded as more valuable than an exact replica simply because it belonged to Diana, which is clearly an extrinsic property of the object. Even though the extra value that accrues to the object is due

to its extrinsic properties, however, it is still the case that this dress is (properly) valued for its own sake, and thus valued non-instrumentally.

Given that value of this sort is possible, then it follows that it could well be the case that we value one true belief over another because of its extrinsic features—i.e., that the one true belief, but not the other, was produced by a reliable cognitive trait that is independently valuable. For example, it could be that we value forming a true belief via a reliable cognitive trait more than a mere true belief because the former belief is produced in such a way that it is of credit to us that we believe the truth. There is thus a crucial lacuna in Zagzebski's argument, and hence she needs to do more to motivate the claim that we must reject the machine-product model of belief in order to respond to the value problem.

A different response to the challenge that Zagzebski raises for reliabilism is given by Michael Brady (2006). In defence of reliabilism, Brady appeals to the idea that to be valuable is to be a fitting or appropriate object of positive evaluative attitudes (such as admiration or love) (e.g., Brentano 1969; Chisholm 1986; Wiggins 1987; Gibbard 1990; Scanlon 1998). That one object is more valuable than another is thus to be understood, on this view, in terms of the fact that that object is more worthy of positive evaluation. Thus, the value problem for reliabilism on this conception of value comes down to the question why knowledge is more worthy of positive evaluation on this view than mere true belief. Brady's contention is that, at least within this axiological framework, it *is* possible for the reliabilist to offer a compelling story about why reliable true belief (and thus knowledge) is more valuable than mere true belief.

Central to Brady's argument is his claim that there are many ways one can positively evaluate something, and thus many different ways something can be valuable. Moreover, Brady argues that we can distinguish *active* from *passive* evaluative attributes, where the former class of attitudes involve pursuit of the good in question. For example, one might actively value the truth, where this involves, for instance, a striving to discover the truth. In contrast, one might at other times merely passively value the truth, such as simply respecting or contemplating it.

With this point in mind, Brady's central thesis is that on the reliabilist account knowledge is more valuable than true belief because certain active positive evaluative attitudes are fitting only with regard to the former (i.e., reliable true belief). In particular, given its intrinsic features, reliable true belief is worthy of active love, whereas an active love of unreliable (i.e., accidental) true belief because of its intrinsic features would be entirely inappropriate because there is nothing that we can do to attain unreliable true belief that wouldn't conflict with love of truth.

This is an intriguing proposal, and certainly opens up a possible avenue of defence against Zagzebski's attack on reliabilism since she doesn't even consider the possibility of applying this axiological framework here. One problem that it faces, however—as Pritchard (2007a: §3) points out—is that it is unclear whether we can make sense of the distinction Brady draws between active and passive evaluative attitudes, at least in the epistemic sphere. When Brady talks of passive evaluative attitudes towards the truth he gives examples like contemplating, accepting, embracing, affirming, and respecting. Some of these attitudes are not obviously positive evaluative attitudes, however. Moreover, some of them are not obviously passive either. For example, is to contemplate the truth really to evaluate it *positively*, rather than simply to consider it? Furthermore, in accepting, affirming or embracing the truth, isn't one *actively* positively evaluating the truth? Wouldn't such evaluative attitudes manifest themselves in the kind of practical action that Brady thinks is the mark of active evaluative attitudes? More needs to be said about this distinction before it can do the philosophical work that Brady has in mind.

4. Virtue Epistemology and the Value Problem

So far this discussion has taken it as given that, whatever problems reliabilism faces in this regard, there are epistemological theories available—some form of virtue epistemology, for example—that can deal with them. But not everyone in the contemporary debate accepts this. Perhaps the best known sceptic in this respect is Jonathan Kvanvig (2003), who in effect argues that while virtue epistemology (along with a form of epistemic internalism) can resolve the primary value problem, the real challenge that we need to respond to is that set by

the secondary value problem; and Kvanvig says that there is no solution available to *that*. On this basis, Kvanvig argues that knowledge is no more valuable than anything which falls short of knowledge and concludes that the focus of epistemology should not be on knowledge at all, but rather on *understanding*, an epistemic standing that Kvanvig thinks is clearly of more value than knowledge and anything which falls short of knowledge.

What Kvanvig says about understanding will be considered below. First though, let us consider the specific challenge that he poses for virtue epistemology. In essence, Kvanvig's argument rests on the assumption that it is essential to any virtue-theoretic account of knowledge—and any internalist account of knowledge as well, for that matter (i.e., an account that makes a subjective justification condition necessary for knowledge possession)—that it also includes an anti-Gettier condition. If this is right, then it follows that even if virtue epistemology has an answer to the primary value problem (and Kvanvig concedes that it does), it will not thereby have an answer to the secondary value problem since knowledge is not simply virtuous true belief. Moreover, Kvanvig argues that once we recognise what a gerrymandered notion a non-Gettierized account of knowledge is, it becomes apparent that there is nothing valuable about the anti-Gettier condition on knowledge that needs to be imposed. But if that is right, then it follows by even virtue epistemic lights that knowledge—i.e., non-Gettierized virtuous true believing—is no more valuable than one of its proper sub-sets—i.e., mere virtuous true believing.

There are at least two aspects of Kvanvig's argument that are problematic. To begin with, it isn't at all clear why an analysis of knowledge should make it obvious why knowledge is valuable, something that seems to be being assumed here. Why should the complexity and awkwardness of a Gettier-proof account of knowledge be any indication that knowledge lacks value (or, conversely, that no adequate analysis has been given of this obviously valuable notion)? Nevertheless, this sort of assumption is common to much of recent epistemological debate, cropping up in the work of virtue epistemologists—see, for example, Zagzebski (1999)—and non-virtue epistemologists alike—see, for example, Timothy Williamson (2000: chapter 1). (For critical discussion of this assumption, see DePaul *forthcoming*).

A more serious problem is that Kvanvig seems not to have noticed that many virtue epistemologists—among them Sosa (1988; 1991; 2007), Zagzebski (e.g., 1996; 1999) and Greco (2002; 2007; *forthcominga*; *forthcomingb*)—think that their view *can* deal with Gettier problems without needing to add an additional anti-Gettier condition on knowledge. The way this is achieved is by making the move noted above of treating knowledge as a state that includes both the truly believing and the virtuous source by which that true belief was acquired. Greco (*forthcominga*: 14), for example, makes a distinction between (i) a belief's being true and virtuously formed, and (ii) a belief's being true *because* virtuously formed. On the virtue-theoretic account he proposes, knowledge is to be analysed as the latter, and it is only when so analysed, claims Greco, that virtue epistemology can respond to the Gettier problem. Kvanvig, however, resolutely reads virtue epistemologists as uniformly offering an account of knowledge cast along the lines of (i).

Moreover, building on earlier work by Sosa and Zagzebski on this score, Greco (e.g., *forthcominga*) argues that it is only if the virtue epistemological proposal is read in this way that it is able to answer the value problem. More specifically (though he does not put the point in these terms), the answer to the value problem offered by virtue epistemology on this construal is able to respond to not only the secondary value problem but also the tertiary value problem. This is because knowledge, on this view, is simply the cognitive aspect of a more general notion, that of achievement. That is, Greco argues that achievements are successes that are because of ability, and thus, given that on his view knowledge is cognitive success (i.e., true belief) that is because of cognitive ability, knowledge is cognitive achievement. The import of this claim to our current discussion is that achievements are, plausibly at least, of final value (Greco himself describes the kind of value in play here as intrinsic value, but it is clear from how he describes it that it is specifically final value that he has in mind). If this is right, then cognitive achievements—i.e., knowledge—will also have final value, and thus one is well on one's way to answering the tertiary value problem (and thus the secondary value problem also).

There are thus two key theses to this account of the value of knowledge—that achievements are finally valuable, and that knowledge is a form of achievement—both of which could be

called into question. As regards the first thesis, one might object that some successes that are because of ability—i.e., achievements, on this view—are too trivial or easy or wicked to count as finally valuable. As Pritchard (2007*b*) argues, however, this line of objection is far from decisive. After all, it is open to Greco to argue that the claim is only that all achievements *qua* achievements are finally valuable, not that the overall value of every achievements is particularly high. It is thus consistent with the proposal that some achievements have a very low—perhaps even negative, if that is possible—value in virtue of their other properties (e.g., their triviality). Indeed, a second option in this regard is to allow that not all achievements enjoy final value whilst nevertheless maintaining that it is in the nature of achievements to have such value (e.g., much in the way that one might argue that it is in the nature of pleasure to be a good, even though some pleasures are bad). Since, as noted above, all that is required to meet the (tertiary) value problem is to show that knowledge is generally distinctively valuable, this claim would almost certainly suffice for Greco's purposes.

In any case, it is the second claim that Greco makes—i.e., that knowledge is to be understood as a kind of achievement—that is the most controversial. There are two key problems with this claim. The first is that there sometimes seems to be more to knowledge than a cognitive achievement; the second is that there sometimes seems to be less to knowledge than a cognitive achievement.

As regards the first claim, notice that achievements seem to be compatible with at least one kind of luck. Pritchard (2007*b*) offers the following example to illustrate this point. Suppose that an archer hits a target by employing her relevant archery abilities, but that the success is 'gettierized' by luck intervening between the archer's firing of the arrow and the hitting of the target. For example, suppose that a freak gust of wind blows the arrow off-course, but then a second freak gust of wind happens to blow it back on course again. The archer's success is thus lucky in the sense that it could very easily have been a failure. When it comes to 'intervening' luck of this sort, Greco's account of achievements is able to offer a good explanation of why the success in question does not constitute an achievement. After all, we would not say that the success was because of the archer's ability in this case.

Notice, however, that not all forms of luck are of this intervening sort. For suppose that nothing intervenes between the archer's firing of the arrow and the hitting of the target. However, the success is still lucky in the relevant sense because, unbeknownst to the archer, she just happened to fire at the only target on the range that did not contain a forcefield which would have repelled the arrow. Is the archer's success still an achievement? Intuition would seem to dictate that it is; it certainly seems to be a success that is because of ability, even despite the luckiness of that success. Achievements, then, are, it seems, compatible with luck of this 'environmental' form even though they are not compatible with luck of the standard 'intervening' form.

The significance of this conclusion for our purposes is that knowledge is incompatible with *both* forms of luck. In order to see this, one only needs to note that an epistemological analogue of the archer case just given is the famous barn façade example. In this example, we have an agent who forms a true belief that there is a barn in front of him. Moreover, his belief is not subject to the kind of 'intervening' luck just noted and which is a standard feature of Gettier-style cases (it is not as if, for example, he is looking at what appears to be a barn but which is not in fact a barn, but that his belief is true nonetheless because there is a barn behind the barn shaped object that he is looking at). Nevertheless, his belief is subject to environmental luck in that he is, unbeknownst to him, in barn façade county in which every other barn-shaped object is a barn façade. Thus, his belief is only luckily true in that he could very easily have been mistaken in this respect. Given that this example is structurally equivalent to the 'archer' case just given, it seems that just as we treat the archer as exhibiting an achievement in that case, so we should treat this agent as exhibiting a cognitive achievement here. The problem, however, is that it is almost universally accepted that the agent in the barn façade case lacks knowledge. Knowledge, it seems, is incompatible with environmental luck in a way that achievements, and thus cognitive achievements, are not.

Greco (2007; *forthcominga*; *forthcomingb*; cf. Pritchard 2007*b*; 2007*c*; *forthcominga*) has made a number of salient points regarding this case. For example, he has argued for a conception of what counts as a cognitive ability according to which the agent in the barn façade case would not count as exhibiting the relevant cognitive ability. Moreover, he has argued that, in any case, there are grounds to think that there may be something special about

the concept of knowledge which would mean that knowledge might be more resistant to certain kinds of luck than achievements more generally.

Even if these claims can be made to stick, however, there is a second problem on the horizon, which is that it seems that there are some cases of knowledge which are not cases of cognitive achievement. One such case is offered by Jennifer Lackey (2007), albeit to illustrate a slightly different point. Lackey asks us to imagine someone arriving at the train station in Chicago who, wishing to obtain directions to the Sears Tower, approaches the first adult passer-by she sees. Suppose the person she asks is indeed knowledgeable about the area and gives her the directions that she requires. Intuitively, any true belief that the agent forms on this basis would ordinarily be counted as knowledge. Indeed, if one could not gain testimonial knowledge in this way, then it seems that we know an awful lot less than we think we know. What is significant about this case, however, is that we would not intuitively regard the truth of the agent's belief as being because of her cognitive abilities. Indeed, if anything, we would think that her cognitive success was down to her *informant's* cognitive abilities. Thus, it seems that there are cases of knowledge which are not also cases of cognitive achievement.

It is worth being clear about the nature of this objection. Lackey takes cases like this to demonstrate that one can possess knowledge without it being of any credit to one that one's belief is true. As Pritchard (e.g., 2007*b*) points out, however, this conclusion is surely too strong, in that the agent *is* employing her cognitive abilities to some degree, and so surely deserves *some* credit for the truth of the belief formed (she would not have asked just anyone, for example, nor would she have simply accepted just any answer given by her informant). The point is thus rather that whatever credit the agent is due for having a true belief, it is not the kind of credit that reflects a *bona fide* cognitive achievement because of how this cognitive success involves 'piggy-backing' on the cognitive efforts of others.

The obvious way in which someone like Greco would respond to this sort of case is to either claim that, despite first appearances, the agent concerned does not have knowledge or else claim that she does have knowledge but that, on closer inspection, this is a genuine cognitive achievement after all. Neither alternative looks particularly appealing, though no doubt a

strong case can be constructed in support of at least one of these responses. The fundamental problem facing the view, however, as Pritchard (2007*b*) points out, is that once one combines this problem with the one mentioned earlier (i.e., the problem that some cognitive achievements don't seem to be cases of knowledge), the view starts to look far from compelling. Indeed, Pritchard (2007*b*) argues that the moral that should be drawn is that there is no adequate response to the tertiary value problem available after all (the virtue-theoretic account being the most promising line on this score) and that a revisionary response to this problem should therefore be favoured. In particular, the claim is that given the close connection between knowledge and cognitive achievements, and the general final value of achievements, it is not surprising that knowledge is thought to be distinctively (i.e., finally) valuable even though closer reflection reveals that it is actually a distinct epistemic standing, that of cognitive achievement, that is distinctively valuable.

5. Understanding and Epistemic Value

As noted above, the main conclusion that Kvanvig (2003) draws from his reflections on the value problem is that the real focus in epistemology should not be on knowledge at all but on understanding, an epistemic standing that Kvanvig does think is especially valuable but which, he argues, is distinct from knowing (i.e., one can have knowledge without the corresponding understanding, and one can have understanding without the corresponding knowledge). It is perhaps this aspect of Kvanvig's book that has prompted the most critical response, so it is worth dwelling on the debate regarding his claims in this respect in a little more detail here.

To begin with, one needs to get clear what Kvanvig has in mind when he talks of understanding, since many commentators have found the conception of understanding that he targets problematic. The two usages of the term 'understanding' in ordinary language that Kvanvig focuses on—and which he regards as being especially important to epistemology—are “when understanding is claimed for some object, such as some subject matter, and when it involves understanding that something is the case.” (Kvanvig 2003: 189) The first kind of understanding he calls “objectual understanding”, the second kind “propositional understanding”. In both cases, understanding requires that one successfully

grasp how one's beliefs in the relevant propositions cohere with other propositions one believes (e.g., Kvanvig 2003: 192, 197-8). This requirement entails that understanding is in a sense factive, and this aspect of Kvanvig's account of understanding has proved especially contentious.

More precisely, this requirement entails that propositional understanding is directly factive, while objectual understanding is at least indirectly factive. I take it that the factivity of propositional understanding on this conception is straightforward, even though the factivity of objectual understanding is not. Here is how Kvanvig explains the indirect nature of the factivity of objectual understanding:

Objectual understanding is, of course, not straightforwardly factive, for only propositions can be true or false. Still, the uses I wish to focus on are ones in which factivity is in the background. For example, to understand politics is to have beliefs about it, and for this objectual understanding to be the kind of interest here requires that those beliefs be true. (Kvanvig 2003: 191)

So even though objectual understanding is not directly factive, it is meant to be at least indirectly factive in the way just outlined. Later on in the book, however, Kvanvig (e.g., 2003: 201-2) seems to weaken this requirement that objectual understanding be factive, allowing that one can have some false beliefs about the subject matter yet still count as having objectual understanding of it (as long as those false beliefs are suitably peripheral to the subject matter).

Given that understanding—propositional understanding at any rate—is factive, Kvanvig's argument for why understanding is distinct from knowledge does not relate to this condition (as we will see in a moment, it is standard to argue that understanding is distinct from knowledge precisely because only understanding is non-factive). Instead, Kvanvig notes two key differences between understanding and knowledge: that understanding, unlike knowledge, admits of degrees, and that understanding, unlike knowledge, is compatible with epistemic luck. Most commentators, however, have tended to focus not on these two theses concerning the different properties of knowledge and understanding, but rather on Kvanvig's

claim that understanding is (at least indirectly) factive.

For example, Elgin (*forthcoming*; cf. Elgin 1996; 2004) and Riggs (*forthcoming*) both argue that it is possible for an agent to have understanding and yet lack true beliefs in the relevant propositions. Central to Riggs's critique is the fact that understanding admits of degrees, something that Kvanvig freely admits. Hence, depending on how well one's beliefs about such a subject matter are integrated into a coherent system, one could count as having a greater or lesser understanding of that subject matter. According to Riggs, however, if that is right then we should immediately be suspicious of the claim that understanding is factive, especially when it comes to objectual understanding. Perhaps a fair number of one's beliefs about a certain subject matter—quantum physics, say—could be false yet one could still count as understanding that subject matter because enough of them are true and there is a great deal of cohesive integration amongst them.

Elgin's (*forthcoming*) challenge builds on Riggs's remarks by looking at the non-factive way in which we talk of understanding in the sciences. She cites a number of cases in which science has progressed from one theory to a better theory where, we would say, understanding has increased in the process even though the theories are, strictly speaking at least, *false*. For example, Kepler's theory of how the Earth travels around the sun is an improvement on that of Copernicus (which was in turn an improvement on the Ptolemaic theory), since it recognises that the orbit of the Earth is elliptical rather than circular. Nevertheless, this theory is regarded as strictly speaking false. Still, didn't our understanding of the heavens increase as we moved from the one theory to the next? A different kind of case that Elgin offers concerns scientific idealisations, such as the ideal gas law. Scientists know full well that no actual gas behaves in this way, yet the introduction of this useful fiction clearly improved our understanding of the behaviour of actual gasses.

It is not clear just how devastating these sorts of criticisms of Kvanvig's account of understanding are. To begin with, they only seem to affect his treatment of objectual understanding, not propositional understanding. Moreover, one could plausibly argue that all they show is that Kvanvig's later concession that objectual understanding need not be entirely factive should be made essential to the view, rather than an afterthought. With that alteration

to the view in mind, the obvious question to ask is whether one can really make sense of a case in which an agent had a large body of false beliefs about a subject matter yet had understanding of that subject matter nonetheless. It is true, of course, that we sometimes talk of someone who has many false beliefs about a subject matter having some limited understanding of the topic, but the way we qualify our ascription of understanding here seems to be an indication that understanding is not actually possessed. Someone with a highly cohesive set of beliefs about a subject matter can, despite many of those beliefs' being false, get closer to having understanding of that subject matter by gaining more true beliefs. But one can grant this point without thereby granting that before and after the agent gained these true beliefs she understood the subject matter. In short, that understanding admits of degrees does not entail that there is no lower limit to what is required for understanding—one can surely get closer to having objectual understanding while nevertheless lacking it.

Elgin's counterexamples from the field of science might be thought to be more persuasive in this regard, though even this is not so clear on closer inspection. Take first the case of Kepler's advance on Copernicus's theory of the heavens. Could not Kvanvig account for what is going on here by saying that while the theory taken as a whole is strictly speaking false, there is nevertheless a great deal of truth within it (e.g., true propositions about the shape of the Earth's orbit)?

Moreover, as just noted, it is entirely consistent with what Kvanvig says about understanding that we might get closer to having understanding of a subject matter as we move from one theory to the next even though neither theory gives us understanding of the subject matter in hand. That is, while the first point could allow Kvanvig to contend that Kepler's theory really is a case of understanding, even though Kepler had some false beliefs about the subject matter in hand, the second point can allow Kvanvig to claim that there is no problem in the idea of there being genuine development towards an understanding of the heavens well before Kepler came on to the scene, even if no-one had enough true beliefs about this subject matter to be truly said to understand. With these two claims in hand, the historical example of the development of our understanding of the relative movements of the Earth and the sun does not seem to pose much of a threat to Kvanvig's view.

One might think that Elgin's other example—that of the scientist's use of idealisations, such as the ideal gas law—poses a more intractable problem for Kvanvig's position, because in this case the scientist has *no* true beliefs at all about the ideal gas law. There are at least two problems with this suggestion. The first is that it would be odd to construe the subject matter in this case as being so narrow when one could more plausibly regard it as taking in a whole range of claims about the behaviour of gases. If this is right, then it isn't at all clear that what we have here is not simply at best a case in which the agent has only a few false beliefs about the subject matter. And remember that on the more refined reading of Kvanvig's account of objectual understanding urged above, one's beliefs can diverge from the truth to a small degree without that entailing that one lacks understanding.

The second reason why this suggestion is problematic is that it is far from obvious that the scientists in this case really do have false beliefs in the subject matter in the first place. Given that scientists recognise—if only implicitly—that the ideal gas law is a useful fiction, why should we suppose that they believe it to be true? And if they don't have any false beliefs in this regard, then there is not even a challenge here to the unrefined conception of objectual understanding that we initially saw Kvanvig proposing.

Indeed, this point has ramifications for the other sort of case that Elgin describes—that of scientific progress from one strictly speaking false theory to another. After all, it is plausible to contend that, at least when it comes to the most highly theoretical aspects of science at any rate, scientists are fully aware that their theories, taken as a whole, are not strictly speaking true (or, at least, that they lack grounds for thinking that they are strictly speaking true). If that's right, then it would be odd to regard the fully rational scientist as actually believing the theory as such, rather than merely believing certain aspects of the theory (e.g., claims about what observations have been made, or the logical relations between different aspects of the theory). That is, scientists might well *accept* their theories in such cases (i.e., endorse them for all practical purposes, as the best theory available), even though they don't actually believe them (cf. van Fraassen 1980). If this is correct, however, then it follows that scientists have far fewer false beliefs than Elgin imagines, and thus the scope for her to show that the case of science reveals that at least scientific understanding is non-factive is severely reduced.

A very different sort of challenge to Kvanvig's treatment of understanding comes from Brogaard (2005). She argues that Kvanvig's claim that understanding is of greater value than knowledge is only achieved because he fails to give a rich enough account of knowledge. More specifically, Brogaard claims that we can distinguish between objectual and propositional knowledge just as we can distinguish between objectual and propositional understanding. Propositional understanding, argues Brogaard, no more requires coherence in one's beliefs than propositional knowledge, and so the difference in value between the two cannot lie here. Moreover, while Brogaard grants that objectual understanding does incorporate a coherence requirement, this again fails to mark a value-relevant distinction between knowledge and understanding because the relevant counterpart—objectual knowledge (i.e., knowledge of a subject matter)—also incorporates a coherence requirement. So provided that we are consistent in our comparisons of objectual and propositional understanding on the one hand, and objectual and propositional knowledge on the other, Kvanvig fails to make a sound case for thinking that understanding is of greater value than knowledge.

A second line of criticism against Kvanvig concerns his claim that knowledge is distinct from understanding on the grounds that only the latter admits of degrees and is compatible with epistemic luck (to keep matters simple, we will focus on propositional understanding in what follows). As Pritchard (*forthcomingb*) notes, the import of the former claim is moot since it does not follow from this claim that there are cases in which knowledge is possessed and yet the corresponding understanding is not, or that there are cases in which understanding is possessed but the corresponding knowledge is not. This point becomes especially important once one notices that the relationship between understanding and epistemic luck may well not be quite as Kvanvig supposes.

Stephen Grimm (2006), for example, argues that understanding is just as incompatible with epistemic luck as knowledge is. In contrast, Pritchard (*forthcomingb*) argues that *both* Grimm and Kvanvig are wrong on this score, in that while understanding is compatible with a certain kind of epistemic luck that knowledge is not compatible with, it is incompatible with a second kind of epistemic luck that knowledge is incompatible with. In particular, Pritchard's

claim is that while understanding is incompatible with the Gettier-style intervening epistemic luck noted earlier, it is compatible with environmental luck.

In order to see what this point involves, consider someone arriving home to find their house in flames and a number of people at the front of the house dressed as fire officers. Suppose that our hero goes up to one of these people and asks why her house burned down and is told that the reason is faulty wiring. As it happens, this is indeed the reason why the house burned down, but it is just a matter of luck that this agent has a true belief in this regard since the person she asked was not a fire officer at all, but merely someone on their way to a fancy dress party dressed up as a fire officer. Intuitively, one cannot come to know why one's house burned down by forming a true belief in the target proposition in this fashion (i.e., one cannot come to know that one's house burned down because of faulty wiring). Equally, however, one cannot gain an understanding of why one's house burned down in this fashion either, and thus, it seems, Kvanvig is wrong to think that understanding is compatible with epistemic luck.

Notice, however, that the type of epistemic luck in play in this example is specifically of the intervening type found in standard Gettier-style cases. Moreover, it is interesting that when we change the example to make the type of epistemic luck in play of the environmental kind then we generate very different intuitions.

Suppose, for example, that our hero finds out why her house burned down by speaking to a genuine fire officer who has first-hand knowledge in this regard. Nevertheless, the belief that our hero forms on this basis is still only luckily true because she could very easily have asked one of the other people outside her house who are dressed as fire officers and yet the person she spoke to is in fact the only genuine fire officer there (as before, the rest are on their way to a fancy dress party). Clearly, the environmental epistemic luck involved here prevents our agent from having knowledge of why her house burned down (i.e., of knowing that her house burned down because of faulty wiring). Nevertheless, it does seem right that one can gain understanding of why one's house burned down in this fashion. After all, the person one spoke to is a genuine fire officer. If this is right, then it seems that while understanding is incompatible with one type of knowledge-undermining epistemic luck (i.e., intervening

epistemic luck), it is entirely compatible with another type of knowledge-undermining epistemic luck (i.e., environmental epistemic luck).

Pritchard's diagnosis for why this might be the case is that understanding, unlike knowledge, is a form of cognitive achievement, for recall we saw him arguing earlier that cognitive achievements, unlike knowledge, are compatible with environmental epistemic luck as well. If this is right, then it enables Pritchard to offer a more explicit account of why understanding is distinctively valuable in the way that Kvanvig alleges. After all, if all achievements are finally valuable, and understanding is itself a type of cognitive achievement, then understanding will be finally valuable too, unlike knowledge on this view.

6. Other Accounts of the Value of Knowledge

John Hawthorne (2004) has recently argued that knowledge is valuable because of the role it plays in practical reasoning. More specifically, Hawthorne (2004: 30) argues for the principle that one should use a proposition p as a premise in one's practical reasoning only if one knows p . Hawthorne primarily motivates this line of argument by appeal to the lottery case. This concerns an agent's true belief that she holds the losing ticket for a fair lottery with long odds and a large cash prize, a belief that is based solely on the fact that she has reflected on the odds involved. Intuitively, we would say that such an agent lacks knowledge of what she believes, even though her belief is true and even though her justification for what she believes—assessed in terms of the likelihood, given this justification, of her being right—is unusually strong. Moreover, were this agent to use this belief as a premise in her practical reasoning, and so infer that she should throw the ticket away without checking the lottery results in the paper for example, then we would regard her reasoning as problematic.

Lottery cases therefore seem to show that justified true belief, no matter how strong the degree of justification, is not enough for acceptable practical reasoning—instead, knowledge is required. Moreover, notice that we can alter the example slightly so that the agent does possess knowledge while at the same time having a *weaker* justification for what she believes (where strength of justification is again assessed in terms of the likelihood, given this justification, that the agent's belief is true). If the agent had formed her true belief by reading

the results in a reliable newspaper, for example, then she would count as knowing the target proposition and can then infer that she should throw the ticket away without criticism. It is more likely, however, that the newspaper has printed the result wrongly than that she should win the lottery. This sort of consideration seems to show that knowledge, even when accompanied by a relatively weak justification, is better (at least when it comes to practical reasoning) than a true belief that is supported by a relatively strong justification but does not amount to knowledge. If this is the right way to think about the connection between knowledge possession and practical reasoning, then it seems to offer at a potential response to at least the secondary value problem.

In response to Hawthorne's claim about the pivotal role of knowledge in practical reasoning, Matthew Weiner (*forthcoming*) has argued that knowledge is not important in itself for practical reasoning. More specifically, he argues that knowledge is what he calls a "Swiss Army Concept", in the sense that when we ascribe knowledge we thereby ascribe several valuable sub-concepts—Weiner lists truth, justification, persistence, stability of justification, and safety (i.e., that one could not have easily been wrong). Each of these sub-concepts could be valuable to us, depending upon which standpoint on our practical reasoning we take, but on no standpoint is knowledge of particular value. Thus, claims Weiner, the value of knowledge relates to how ascribing knowledge is a shorthand way of ascribing a number of valuable sub-concepts, each of which may be of particular value in our practical reasoning depending upon what standpoint on our practical reasoning we take. But knowledge has no special value in itself, at least as regards practical reasoning.

For example, one standpoint that could be taken on your practical reasoning is to care about whether things turn out well for you. From this standpoint, argues Weiner, it is truth that is especially important. Consider again a case in which one reasons from one's true belief that one owns a losing lottery ticket to the conclusion that one should not bother checking the lottery result in the paper. Although this reasoning might seem generally suspicious, as we noted above, from this specific perspective it is unobjectionable—after all, things do turn out well for you in this case because, since your premises are true, you rightly save yourself the trouble of finding out what the local newspaper says about the lottery result.

In contrast, from other standpoints, such as that of caring that your reasoning not be vulnerable to criticism, a different result will be generated. (In the case under consideration, for example, it would be remiss not to check the lottery result in a local newspaper given that there is a chance that you have won a large cash prize, and given also that making such a check would not be unduly onerous). From different perspectives, then, the very same practical inference could be assessed differently. But from no perspective, argues Weiner, does it matter that one knows one's premises.

Whether or not Weiner is right about this, it is an intriguing possibility that the value placed on knowledge might derive from its being a Swiss Army Concept in this way. For it highlights the fact that how we understand the concept of knowledge can have important ramifications for how we go about determining the special value, if any, of knowledge.

A second author who thinks that our understanding of the concept of knowledge can have important ramifications for the value of knowledge is Edward Craig (1990). Craig's project begins with a thesis about the value of the concept of knowledge. Simplifying somewhat, Craig hypothesises that the concept of knowledge is important to us because it fulfils the valuable function of enabling us to identify reliable informants. The idea is that it is obviously of immense practical importance to be able to recognise those from whom we can gain true beliefs, and that it was in response to this need that the concept of knowledge arose. As with Hawthorne's theory, this proposal, if correct, could potentially offer a resolution of at least the secondary value problem.

What is particularly interesting about Craig's approach for our present purposes is that he claims that the concept of knowledge has evolved over time away from its original function through a process Craig calls "objectification". In essence, the process of objectification occurs because the need to eliminate error that is built-into the concept of knowledge becomes 'stretched' to accommodate increasingly demanding error-possibilities as we become intellectually more sophisticated. This is why, according to Craig, we have ended up with a concept of knowledge that sometimes denies knowledge to those who are clearly good informants (e.g., when some far-fetched error-possibility is made salient), even though the original function was to enable us to identify reliable informants.

This proposal that the concept of knowledge may have changed over time so that what we now call ‘knowledge’ may sometimes perform a different function to the one that our original concept of knowledge was supposed to track is obviously of central importance to debates about the value of knowledge, as Craig’s account of objectification indicates. After all, if we make the plausible assumption, with Craig, that at least the original function of knowledge was to pick out some property of agents that was valuable to us, then on this picture of an ‘evolving’ concept we can both account for the fact that our present-day conception of knowledge seems to be of particular value even though there are cases in which knowledge is ascribed where it is not of any distinctive value. If this is right, then we should not look to the necessary and sufficient conditions for knowledge in order to determine why it seems to be of distinctive value to us, but rather merely to its original function.

7. The Value of True Belief

So far, in common with most of the contemporary literature in this regard, we have tended to focus on the value of knowledge relative to other epistemic standings. A related debate in this respect, however—one that has often taken place largely in tandem with the mainstream debate on the value of knowledge—has specifically concerned itself with the value of true belief and I will close by considering this issue.

Few commentators treat truth or belief as being by themselves valuable (though see Kvanvig 2003: ch. 1), but it is common to treat true belief as valuable, at least instrumentally. True beliefs are clearly often of great practical use to us. The crucial *caveat* here, of course, concerns the use of the word ‘often’. After all, it is also often the case that a true belief might actually militate against one achieving one’s goals, as when one is unable to summon the courage to jump a ravine and thereby get to safety, because one knows that there is a serious possibility that one might fail to reach the other side. In such cases it seems that a false belief in one’s abilities would be better than a true belief, if the goal in question (jumping the ravine) is to be achieved.

Moreover, some true beliefs are beliefs in trivial matters, and in these cases it isn’t at all clear

why we should value such beliefs at all. Imagine someone who, for no good reason, concerns herself with measuring each grain of sand on a beach, or someone who, even while being unable to operate a telephone, concerns herself with remembering every entry in a foreign phonebook. Such a person would thereby gain lots of true beliefs but, crucially, one would regard such truth-gaining activity as rather pointless. After all, these true beliefs do not obviously serve any valuable purpose, and so do not seem to have any instrumental value (or, at the very least, what instrumental value these beliefs have is vanishingly small). It would, perhaps, be better—and thus of greater value—to have fewer true beliefs, and possibly more false ones, if this meant that the true beliefs that one had concerned matters of real consequence.

At most, then, we can say that true beliefs often have instrumental value. What about final (or intrinsic) value? One might think that if the general instrumental value of true belief was moot then so too would be the intuitively stronger thesis that true belief is generally finally valuable. Nevertheless, many have argued for such a claim.

One condition that seems to speak in favour of this thesis is that as truth seekers we are naturally curious about what the truth is, even when that truth is of no obvious practical import. Accordingly, it could be argued that from a purely epistemic point of view, we do regard all true belief as valuable for its own sake, regardless of what further prudential goals we might have (e.g., Goldman 1999: 3; Lynch 2004: 15-16; Alston 2005: 31). Curiosity will only take you so far in this regard, however, since we are only curious about certain truths, not all of them. To return to the examples given a moment ago, no fully rational agent is curious about the measurements of every grain of sand on a given beach, or the name of every person in a random phonebook (i.e., no rational person wants to know these truths independently of having some prudential reason for knowing them).

Still, one could argue for a weaker claim and merely say that it is *prima facie* or *pro tanto* finally good to believe the truth (cf. David 2005; Lynch *forthcoming*), where cases of trivial truths such as those just given are simply cases where, *all things considered*, it is not good to believe the truth. After all, we are familiar with the fact that something can be *prima facie* or *pro tanto* finally good without being all-things-considered good. For example, it may be

finally good to help the poor and needy, but not all-things-considered good given that helping the poor and needy would prevent you from doing something else which is at present more important (such as saving that child from drowning).

At this point one might wonder why it matters so much to (some) epistemologists that true belief is finally valuable. Why not instead just treat true belief as often of instrumental value and leave the matter at that? The answer to this question lies in the fact that many want to regard truth—and thereby true belief—as being in some sense the fundamental epistemic goal. Accordingly, if true belief is not finally valuable—and only typically instrumentally valuable—then this seems to downplay the status of the epistemological project.

There are a range of options here. The conservative option is to contend that truth is the fundamental goal of epistemology and also contend that true belief is finally valuable—at least in some restricted fashion. Marian David (2001; 2005) falls into this category. In contrast, one might argue that truth is the fundamental goal while at the same time claiming that true belief is *not* finally valuable. Sosa (see especially 2004, but also 2000; 2003) seems (almost) to fall into this camp, since he claims that while truth is the fundamental epistemic value, we can accommodate this thought without having to thereby concede that true belief is finally valuable.

Another axis on which this debate can be configured is in terms of whether one opts for an epistemic-value monism or an epistemic-value pluralism—that is, whether one thinks there is only one fundamental epistemic goal, or several. Kvanvig (e.g., 2005) would be an example of someone who endorses epistemic-value pluralism, since he thinks that there are a number of fundamental epistemic goals, with each of them being of final value. Crucial to Kvanvig's argument is that there are some epistemic goals which are not obviously truth-related—he cites the examples of having an empirically adequate theory, making sense of the course of one's experience, and inquiring responsibly. This is important because if the range of epistemic goals that Kvanvig identified were all truth-related, then it would prompt the natural response that such goals are valuable only because of their connection to the truth, and hence not fundamental epistemic goals at all.

Presumably, though, it ought also to be possible to make a case for an epistemic-value pluralism where the fundamental epistemic goals were not finally valuable (or, at least, *à la* Sosa, where one avoided taking a stance on this issue). More precisely, if an epistemic-value monism that does not regard the fundamental epistemic goal as finally valuable can be made palatable, then there seems no obvious reason why a parallel view that opted for pluralism in this regard could not similarly be given a plausible supporting story.

Bibliography

Works Cited

- Alston, W. (2005), *Beyond “Justification”*: Dimensions of Epistemic Evaluation, Ithaca, NJ: Cornell University Press.
- Brady, M. S. (2006). ‘Appropriate Attitudes and the Value Problem’, *American Philosophical Quarterly* 43: 91-99.
- Brady, M. S., and Pritchard, D. H., eds. (2003), *Moral and Epistemic Virtues*, Oxford: Blackwell.
- Brentano, F. (1969), *The Origin of Our Knowledge of Right and Wrong*, tr. R. Chisholm, London: Routledge & Kegan Paul.
- Brogaard, B. (2005), ‘I Know. Therefore, I Understand’, *manuscript* [available [here](#)].
- ——— (2007), ‘Can Virtue Reliabilism Explain the Value of Knowledge?’, *Canadian Journal of Philosophy* 36: 335-54.
- Chisholm, R. (1986), *Brentano and Intrinsic Value*, Cambridge: Cambridge University Press.
- Craig, E. (1990), *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Oxford: Clarendon Press.
- David, M. (2001), ‘Truth as the Epistemic Goal’, in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Virtue, and Responsibility*, ed. M. Steup, Oxford: Oxford University Press.
- ——— (2005), ‘Truth as the Primary Epistemic Goal: A Working Hypothesis’, in *Contemporary Debates in Epistemology*, eds. E. Sosa and M. Steup, 296-312, Oxford:

Blackwell.

- DePaul, M. (1988), ‘Is Truth our Epistemic End?’, *unpublished manuscript*.
- ——— (1993). *Balance and Refinement: Beyond Coherence Methods of Moral Inquiry*, New York: Routledge.
- ——— (*Forthcoming*), ‘Does an Ugly Analysis Entail that the Target of the Analysis Lacks Value?’, in *Epistemic Value*, eds. A. Haddock, A. Millar and D. H. Pritchard, Oxford: Oxford University Press.
- Elgin, C. (1996), *Considered Judgement*, Princeton, NJ: Princeton University Press.
- ——— (2004), ‘True Enough’, *Philosophical Issues* 14: 113-31.
- ——— (*Forthcoming*), ‘Is Understanding Factive?’, in *Epistemic Value*, eds. A. Haddock, A. Millar and D. H. Pritchard, Oxford: Oxford University Press.
- Gibbard, A. (1990), *Wise Choices, Apt Feelings*, Oxford: Clarendon Press.
- Goldman, A. (1999), *Knowledge in a Social World*, Oxford: Oxford University Press.
- Goldman, A., and Olsson, E. (*Forthcoming*), ‘Reliabilism and the Value of Knowledge’, in *Epistemic Value*, eds. A. Haddock, A. Millar and D. H. Pritchard, Oxford: Oxford University Press.
- Greco, J. (1999), ‘Agent Reliabilism’, *Philosophical Perspectives* 13: 273-96.
- ——— (2000), *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*, Cambridge: Cambridge University Press.
- ——— (2002), ‘Knowledge as Credit for True Belief’, in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, eds. M. DePaul and L. Zagzebski, Oxford: Oxford University Press.
- ——— (2007), ‘The Nature of Ability and the Purpose of Knowledge’, *typescript*.
- ——— (*Forthcominga*), ‘The Value Problem’, in *Epistemic Value*, eds. A. Haddock, A. Millar and D. H. Pritchard, Oxford: Oxford University Press.
- ——— (*Forthcomingb*), ‘What’s Wrong With Contextualism?’, *The Philosophical Quarterly*.
- Grimm, S. (2006), ‘Is Understanding a Species of Knowledge?’, *British Journal for the Philosophy of Science* 57: 515-36.
- Hawthorne, J. (2004), *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Jones, W. (1997), ‘Why Do We Value Knowledge?’, *American Philosophical*

Quarterly 34: 423-40.

- Kaplan, M. (1985), 'It's Not What You Know That Counts', *Journal of Philosophy* 82: 350-63.
- Kvanvig, J. (2003), *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.
- ——— (2005), 'Truth is not the Primary Epistemic Goal', in *Contemporary Debates in Epistemology*, eds. E. Sosa and M. Steup, 285-96, Oxford: Blackwell.
- Lackey, J. (2007), 'Why We Don't Deserve Credit for Everything We Know', *Synthese* 156: ??-??.
- Lynch, M. (2004), *True to Life: Why Truth Matters*, Cambridge, MA: MIT Press.
- ——— (Forthcoming), 'The Values of Truth and the Truth of Values', in *Epistemic Value*, eds. A. Haddock, A. Millar and D. H. Pritchard, Oxford: Oxford University Press.
- Percival, P. (2003), 'The Pursuit of Epistemic Good', *Metaphilosophy* 34: 29-47; and reprinted in Brady and Pritchard (2003).
- Pritchard, D. H. (2007a), 'Recent Work on Epistemic Value', *American Philosophical Quarterly* 44: 85-110.
- ——— (2007b), 'The Value of Knowledge', *manuscript* [available [here](#)].
- ——— (2007c), 'Virtue Epistemology and Epistemic Luck, Revisited', *Metaphilosophy* 38: ??-??.
- ——— (Forthcominga), 'Greco on Knowledge: Virtues, Contexts, Achievements', *The Philosophical Quarterly*.
- ——— (Forthcomingb), 'Knowledge, Understanding and Epistemic Value', *Think* 16; to be reprinted in *Epistemology (Royal Institute of Philosophy Lectures)*, (ed.) A. O'Hear, Cambridge: Cambridge University Press.
- Rabinowicz, W., and Roennow-Rasmussen, T. (1999), 'A Distinction in Value: Intrinsic and For its Own Sake', *Proceedings of the Aristotelian Society* 100: 33-49.
- ——— (2003), 'Tropic of Value', *Philosophy and Phenomenological Research* 66: 389-403.
- Riggs, W. (2002), 'Reliability and the Value of Knowledge', *Philosophy and Phenomenological Research* 64: 79-96.

- ——— (*Forthcoming*), ‘Getting the Meno Requirement Right’, in *Epistemic Value*, eds. A. Haddock, A. Millar and D. H. Pritchard, Oxford: Oxford University Press.
- Scanlon, T. M. (1998), *What We Owe to Each Other*, Cambridge, MA: Harvard University Press.
- Sosa, E. (1985), ‘Knowledge and Intellectual Virtue’, *The Monist* 68: 224-45.
- ——— (1988), ‘Beyond Skepticism, to the Best of our Knowledge’, *Mind* 97: 153-89.
- ——— (1991), *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge, UK: Cambridge University Press.
- ——— (2000), ‘For the Love of Truth’, in *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, eds. A. Fairweather and L. Zagzebski, Oxford: Oxford University Press.
- ——— (2003), ‘The Place of Truth in Epistemology’, in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, eds. L. Zagzebski and M. DePaul, Oxford: Oxford University Press.
- ——— (2004), ‘Reply to Linda Zagzebski’, in *Sosa and His Critics*, ed. J. Greco, Oxford: Blackwell.
- ——— (2007), *Apt Belief and Reflective Knowledge*, Oxford: Oxford University Press.
- Swinburne, R. (1999), *Providence and the Problem of Evil*, Oxford: Oxford University Press.
- ——— (2000), *Epistemic Justification*, Oxford: Oxford University Press.
- van Fraassen, B. (1980), *The Scientific Image*, Oxford: Clarendon Press.
- Weiner, M. (*Forthcoming*), ‘Does Knowledge Matter?’, in *Epistemic Value*, eds. A. Haddock, A. Millar and D. H. Pritchard, Oxford: Oxford University Press.
- Wiggins, D. (1987), ‘A Sensible Subjectivism?’, in his *Needs, Values, Truth*, Oxford: Blackwell.
- Williamson, T. (2000), *Knowledge and Its Limits*, Oxford: Oxford University Press.
- Zagzebski, L. (1996), *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge: Cambridge University Press.
- ——— (1999). ‘What is Knowledge?’, in *Epistemology*, eds. J. Greco and E. Sosa, 92-116, Oxford: Blackwell.
- ——— (2003), ‘The Search for the Source of the Epistemic Good’, *Metaphilosophy* 34:

12-28; and reprinted in Brady and Pritchard (2003).

Collections

- Haddock, A., Millar, A. and Pritchard, D. H., eds. (*Forthcoming*), *Epistemic Value*, Oxford: Oxford University Press.

Other Important Works

- DePaul, M. (2001), 'Value Monism in Epistemology', in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Virtue, and Responsibility*, ed. M. Steup, Oxford: Oxford University Press.
- Ellis, B. (1988), 'Solving the Problem of Induction Using a Values Based Epistemology', *British Journal for the Philosophy of Science* 39: 141-60.
- Fallis, D. (2004), 'Epistemic Value Theory and Information Ethics', *Minds and Machines* 14: 101-17.
- ——— (2005), 'Epistemic Value Theory and Judgment Aggregation', *Episteme* 2: 39-55.
- Horwich, P. (2006), 'The Value of Truth', *Noûs* 40: 347-60.
- Kvanvig, J. (1992), *The Intellectual Virtues and the Life of the Mind: On the Place of the Virtues in Contemporary Epistemology*, Savage, Maryland: Rowman & Littlefield.
- ——— (1998), 'Why Should Inquiring Minds Want to Know?: *Meno* Problems and Epistemological Axiology', *The Monist* 81: 426-51.
- ——— (2004), 'Nozickian Epistemology and the Value of Knowledge', *Philosophical Issues* 14: 201-18.
- Pritchard, D. H. (*Forthcoming*), 'Knowing the Answer, Understanding and Epistemic Value', *Grazer Philosophische Studien*.
- Riggs, W. (1998), 'What are the "Chances" of Being Justified?', *The Monist* 81: 452-72.
- ——— (2002), 'Beyond Truth and Falsehood: The Real Value of Knowing that P', *Philosophical Studies* 107: 87-108.

- — (2003), ‘Understanding Virtue and the Virtue of Understanding’, in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, eds. L. Zagzebski and M. DePaul, Oxford: Oxford University Press.
- — (2007). ‘The Value Turn in Epistemology’, in *New Waves in Epistemology*, eds. V. Hendricks and D. H. Pritchard, London: Palgrave Macmillan.
- Sosa, E. (2000), ‘Skepticism and Contextualism’, *Philosophical Issues* 10: 1-18.
- Zagzebski, L. (2001), ‘Recovering Understanding’, in *Knowledge, Truth, and Obligation: Essays on Epistemic Justification, Virtue, and Responsibility*, ed. M. Steup, Oxford: Oxford University Press.
- — (2003), ‘Intellectual Motivation and the Good of Truth’, in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, eds. M. DePaul and L. Zagzebski, 135-54, Oxford: Oxford University Press.
- — (2004), ‘Epistemic Value Monism’, in *Sosa and His Critics*, ed. J. Greco, Oxford: Blackwell.

Other Internet Resources

- [Epistemic Value](#), a weblog devoted to the topic of epistemic value.

Related Entries

[Virtue Epistemology](#)