

DUNCAN PRITCHARD

RECENT WORK ON EPISTEMIC VALUE

ABSTRACT. Recent discussion in epistemology has seen a huge growth in interest in the topic of epistemic value. In this paper I describe the background to this new movement in epistemology and critically survey the contemporary literature on this topic.

1. BACKGROUND

One of the most interesting developments in recent work in epistemology has been its focus on epistemic value. Indeed, it has been suggested—by Wayne Riggs (2006)—that contemporary epistemology is presently undergoing a “value-turn”, such is the impact that value theory is having on epistemological debate. The aim of this survey piece is to provide an overview of this new movement.

This renewed focus on epistemic value primarily grew out of one of the last big developments in epistemology—that of virtue epistemology, which rose to prominence from the mid 1980s onwards.¹ The essential idea behind virtue epistemology, at least in its most robust guises, was that knowledge should essentially and exclusively understood in terms of the epistemic virtues (where the notion of an epistemic virtue was often understood quite broadly so that it included cognitive faculties). Once epistemologists had started to think about knowledge in these virtue-theoretic terms, however, it was unsurprising that they started to ask value-driven questions about their subject matter.

For one thing, the interest in epistemic virtues stimulated a *rapprochement* between epistemology and ethics, with epistemologists looking to ethical and meta-ethical debates for ideas applicable to epistemological problems. Since the issue of value obviously looms large in ethics, it is natural that such investigations would draw considerations about value into epistemological debates. Moreover, virtue theory brings with it a host of theses about value

anyway. Virtuous character traits are meant to be traits that have a special value (though whether this thesis applies to cognitive faculties is not so clear), and the life led by the virtuous agent is also held to be of special value, possibly as a result of the fact that it involves the operation of the epistemic virtues (this depends on one's virtue theory).

This interest in epistemic value that grew out of work on virtue epistemology neatly dove-tailed with a second trend in epistemology that started to gather pace around the mid 1980s. This wasn't a new epistemological movement as such, but rather a general and growing dissatisfaction with one of the main directions of research in epistemology, what we might loosely call 'post-Gettier epistemology'. Ever since Gettier's seminal article back in the early 1960s which showed that knowledge is not justified true belief (at least given a certain plausible construal of justification), the search had been on for a new analysis of knowledge which was Gettier-proof. Many of the ensuing analyses had been found to be themselves subject to Gettier-style problems, and this prompted a large-scale industry of fine-grained work on this issue, resulting in a number of increasingly complex proposals.

The dissatisfaction with this research programme partly arose out of the fact that it didn't seem to be getting anywhere—an intuitive, substantive, and Gettier-proof analysis of knowledge did not seem to be available, at least insofar as one restricted oneself to the usual epistemological tools (some virtue epistemologists think that appeal to virtues can do the trick, for example). This was not the only reason why epistemologists were troubled by this research programme though, for a secondary concern being voiced here—especially by the early virtue epistemologists—was that these complex analyses of knowledge made it mysterious why knowledge should be the sort of thing that we should care about. That is, a growing consensus emerged that an analysis of knowledge should make it clear why knowledge is the sort of thing that we care about, such that if one's analysis of knowledge did not make this clear then that was a pretty serious strike against it.²

2. THE VALUE PROBLEM

This new work on epistemic value was largely directed at the question of why knowledge was distinctively valuable (especially in contrast to other lesser epistemic standings, like merely truly believing). This question has come to be known as the *value problem*. In this

survey I will follow suit with the contemporary literature and take the value problem as my main point of departure (though, as we will see, this question naturally leads on to other issues regarding epistemic value, such as whether there are epistemic standings of more value than knowledge which should accordingly occupy a more central place in epistemology).

The question of why knowledge is distinctively valuable has an important historical precedent in Plato's *Meno*, when Socrates puts 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 equal practical use as knowledge—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 *primary value 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 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 less prosaically, the point being made here is that knowledge, unlike mere true belief, gives one a confidence that is not easily lost, and that it is this property which 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 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 now 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).

We will return to consider other positive proposals for dealing with the primary value problem below. First, however, I want to distinguish the primary value problem from a second issue raised by those questioning the value of knowledge, what I shall call the *secondary value problem*. This 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 which deals with Gettier-style cases. Suppose further, however, that justification adds value to a mere true belief. If this last point is right, then it follows that we could use the fact that knowledge entails justification as a way of dealing with the primary value problem, since there would now be a property of knowledge which mere true belief lacks 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.

How one understands knowledge could thus have an impact on whether one's response to the primary value problem is also able to deal with the secondary value problem. Unfortunately, however, many commentators in this debate have not spotted this fact and have accordingly simply taken the primary value problem to be *the* value problem. In what follows, we will keep these two problems apart. Unless I specify otherwise, when I talk simply of the value problem, I will have the *primary* value problem in mind.

3. RELIABILISM AND THE VALUE 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. The epistemologists driving this line of argument tended to be virtue epistemologists who regarded their own view as being in a far better position to deal with this challenge (we will return to assess virtue-theoretic responses to the value problem in a moment).

A fairly clear statement of what is at issue here is given in a number of places by Linda Zagzebski (e.g., 1996; 2003a).³ 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 not improved upon by being 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.⁴ 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.

With the problem expressed in this way, however, it seems that it will only affect ‘pure’ reliabilist theories—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 value they possess in virtue of being reliable (i.e., that they have some 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—we will consider Greco's proposal in this regard in a moment).

Zagzebski's diagnosis of what is motivating this problem for reliabilism seems to explicitly exclude such a counter-response, however. 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 2003a), where the product is external to the cause.⁵ It is not clear what Zagzebski means by this point, exactly, but she thinks it shows that even where the reliable process is independently valuable (i.e., independent 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 to a good cup of coffee.

Perhaps the best way of getting a handle on 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 where we explicitly evaluate actions in terms of the motives which 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 by 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 means by which the knowledge was acquired. If, in contrast, we have a conception of knowledge which incorporates the way the knowledge was acquired into the state of knowledge, then Zagzebski argues that we can avoid this problem.

Zagzebski's contention is that once one effects this transition away from the machine-product model of belief, then 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 (2002; *forthcoming*) have argued for a ‘credit’ version of virtue epistemology, such that the agent, in virtue of bringing about the positively valuable outcome of a true belief, is due credit as a result. Rather than treat the extra value of knowledge over true belief as deriving from the agent’s achievement in gaining 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, and 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, and yet nevertheless 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 an intrinsically 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 (*forthcoming*), who claims that an intrinsically valuable cause can indeed confer value on its effect in the relevant cases.⁶ Brogaard claims that virtue epistemologists like Zagzebski and Riggs endorse this claim because they adhere to what she call a “Moorean” conception of value such that 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 to 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 where this value arises because they are extrinsically related to something else that we value.

The standard example in this regard is Princess Diana’s dress, which 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 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 (and intrinsically) 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. Accordingly, all Zagzebski has shown is that reliabilists cannot account for the value of knowledge over true belief simply by appealing to the reliability of the process involved. So long as the reliabilist can show that the relevant reliable processes are intrinsically valuable, then a possible response to the value problem is left open to them. Hence, Zagzebski 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.⁷

A different response to the challenge for reliabilism that Zagzebski raises 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).⁸ That one object is more valuable than other 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 of 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 in this regard is his claim that there are many ways in which one can positively evaluate something, and thus on this view lots of different ways in which something can be valuable. Moreover, Brady argues that we can distinguish between *active* and *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 knowledge is more valuable than true belief on the reliabilist account because certain active positive evaluative attitudes are only fitting 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, is that it is unclear that 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 *positively* evaluate it, 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? I think more needs to be said here about this distinction.

4. VIRTUE EPISTEMOLOGY AND THE VALUE PROBLEM

So far we have taken it as given that, whatever the problems that reliabilism faces in this regard, there are epistemological theories available—virtue epistemology, for example (the right sort of virtue epistemology at any rate)—that can deal with the problem. Not everyone in the contemporary debate accepts this, however. Perhaps the most famous sceptic in this respect is Jonathan Kvanvig (2003), who 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 any proper subset of knowledge and concludes that the focus of epistemology should not be on knowledge at all, but on *understanding*, an epistemic standing which Kvanvig thinks is clearly of more value than true belief, knowledge, and any proper subset of knowledge.

We will return to consider what Kvanvig says about understanding in a moment. First though, I want to set out 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 which makes a subjective justification condition necessary for knowledge possession)—that it also includes an anti-Gettier condition as well. 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 won't thereby have an answer to the secondary value problem since knowledge isn't just 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's 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 at all that knowledge lacks value (or, conversely, that we haven't given an adequate analysis 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).⁹

A more serious problem is that Kvanvig seems not to have noticed that many virtue epistemologists—among them Zagzebski (e.g., 1996; 1999) and Greco (2002; *forthcoming*)—think that their view *can* deal with Gettier problems without needing to add an additional anti-Gettier condition on knowledge.¹⁰ The way this is achieved is by making the move noted above of treating knowledge as a state that includes both truly believing and the virtuous source by which that true belief was acquired. Greco (*forthcoming*, 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 so analysed, claims Greco, that virtue epistemology can respond to the Gettier problem.

Kvanvig, however, resolutely reads virtue epistemology as offering a view in line with (i) rather than (ii), and so treats the account as subject to Gettier-style examples. For example, Kvanvig (2003, 107) argues that virtue-theoretic accounts of knowledge which don't include an anti-Gettier condition will be subject to barn façade cases where the agent virtuously forms a true belief by looking at a real barn, even though she is surrounded by barn façades (and isn't able to tell the difference), and so lacks knowledge. As Greco (*forthcoming*, 16) points out, however, it isn't at clear that such cases pose a problem for virtue epistemology provided that it is understood along the lines suggested in (ii). After all, claims Greco, it is not *because* of the virtuous way that the belief was formed that it is true, but rather because of some incidental feature of the environment (that the agent happened to be looking at the one real barn). If this is right, then Kvanvig's argument against virtue-theoretic responses to the secondary value problem does not go through.

While it is certainly a lacuna in Kvanvig's argument that he fails to take this aspect of (some forms of) virtue epistemology seriously, it is not clear to me that the virtue-theoretic response to Gettier problems is as compelling as commentators like Greco suppose. Take the example just given, for example. Is it really clear here that we wouldn't say that the agent's

belief was true *because of* the operation of her virtues? And if we don't say this in this case, then why do we treat agents as forming true beliefs through their virtues in other cases where there is the possibility of error close-by in the neighbourhood (not all such cases are Gettier-style cases after all)? While there is a question mark having over this issue, there is likewise a possibility that the secondary value problem that Kvanvig identifies for virtue epistemology could be resurrected even against those virtue epistemological theories which try to avoid adding an anti-Gettier condition to their analysis of knowledge.¹¹

5. UNDERSTANDING AND EPISTEMIC VALUE

As I noted a moment ago, 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 understanding, an epistemic standing that Kvanvig does think is especially valuable. It is perhaps this aspect of Kvanvig's book that has prompted the most critical response, and so it is worth dwelling on the debate regarding Kvanvig's claims in this respect in a little more detail here.

To begin with, we need 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 grasps how one's belief(s) in the relevant proposition(s) coheres with other propositions which one believes (see, for example, 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 and yet still count as having objectual understanding of it (as long as, at any rate, the false beliefs in question are suitably peripheral to the subject matter). Nevertheless, even though objectual understanding is not directly factive—and may not even be *entirely* factive—many commentators have found the factivity of both these types of understanding to be implausible.

In particular, several commentators have argued that understanding is often not factive, even when we restrict our attentions to propositional or objectual understanding in the way that Kvanvig proposes. For example, Elgin (*forthcoming*) and Riggs (*forthcoming*) each offer cases in which an agent would be plausibly said to have understanding and yet lacks true beliefs in the relevant propositions. Central to Rigg's critique is the fact that understanding admits of degrees, something that Kvanvig freely admits. Hence, depending on how well my beliefs about such a subject matter are integrated into a coherent system, I could count as having a greater or less understanding of that subject matter. According to Riggs, however, if that's right then we should be immediately suspicious of the claim that understanding is factive, especially when it comes to objectual understanding. Perhaps a fair number of my beliefs about a certain subject matter—quantum physics, say—could be false and yet I still count as understanding that subject matter because enough of them are true and there is a great deal of cohesive integration amongst my beliefs in this regard.

Elgin's 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 Copernicus's (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 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, and yet the introduction of this useful fiction clearly improved our understanding of the behaviour of actual gasses.

It's not clear just how devastating these sorts of criticisms against 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 afterthought. With that alteration to the view in mind, the obvious question to ask is whether we could really make sense of a case in which an agent had a large body of false beliefs about a subject matter and yet had understanding of that subject matter nonetheless. It's 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 I am inclined to treat the way we qualify our ascription of understanding here as being an indication that understanding is not actually possessed. Someone with a highly cohesive set of beliefs about a subject matter—but where many of these beliefs are false—can 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 clear on closer inspection. Take the case of Kepler's advance on Copernicus's theory of the heavens first. Couldn't 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 the Earth's orbit)?

Moreover, as just noted, it is entirely consistent with what Kvanvig says about understanding that understanding might increase as we move from 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 in our 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 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 in 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. 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.¹²

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, then Kvanvig fails to make a sound case for thinking that understanding is of greater value than knowledge.

It is difficult to assess Brogaard's arguments without engaging in a more general discussion of what is required for knowledge, so here is perhaps not the best place to explore these arguments further. Nevertheless, what is clear is that the distinction that Brogaard draws between objectual and propositional knowledge can cast a great deal of light on this discussion. Consider, for example, a second point that Riggs (*forthcoming*) makes regarding Kvanvig's treatment of understanding. Some complications aside, the general conception of understanding that Kvanvig has treats knowledge as a sub-component of understanding, such that understanding is like a kind of refined knowing. Riggs argues against this conception by

claiming that the possibility of non-factive understanding shows that it ought to be possible to have understanding of a subject matter even whilst lacking knowledge.

We saw above that while objectual understanding may be consistent with the agent concerned having some false beliefs about a subject matter—something that Kvanvig, eventually at any rate, grants—it is far from clear that we can make sense of an agent understanding a subject matter while having a large body of false beliefs about that subject matter. With this point in mind, and bearing in mind also the distinction between propositional and objectual knowledge (such that it is clear that we are talking about the latter in this regard), it ought to be obvious that Rigg’s argument rests on the contentious assumption that objectual knowledge, like objectual understanding, is factive across the board. After all, it is only if this is so that the fact that one can have objectual understanding of a subject matter even while having some false beliefs about it will translate into the claim that in such cases one lacks objectual knowledge of that subject matter even whilst having objectual understanding of it. But why should we hold that objectual knowledge requires true beliefs across the board regarding the subject matter in question? Can I not know quantum theory even whilst having some false beliefs in this regard? The trick, it seems, is to remember that objectual knowledge is very different from propositional knowledge. It is certainly true, after all, that in this case you do not have propositional knowledge of all the propositions you believe about quantum theory since some of the propositions you believe are false. Nevertheless, once we remember that what is at issue here is objectual knowledge rather than propositional knowledge, then this sort of consideration fails to have any force.¹³

6. THE VALUE OF KNOWLEDGE

Kvanvig’s book has not been the only account offered of the value of knowledge, still less the only account which treats knowledge as not especially valuable. Mark Kaplan (1985) has argued, for example, the existence of Gettier-style cases indicate just why we shouldn’t care about knowledge specifically. After all, in a Gettier-style case we have both true belief and justification. Why, then, should it matter that, due to some incidental quirk in the environment, one fails to know? In short, Kaplan’s point is that having knowledge, as opposed to mere having justified true belief, is of no practical importance, and hence of no

especial importance at all. Note that whether Kaplan is right on this score may be in part dependent upon whether there are accounts of knowledge available that are able to deal with the Gettier cases in a non *ad hoc* manner (such as the type of virtue epistemology that we saw offered by Zagzebski *et al* above).¹⁴

In contrast to Kaplan, some commentators—most recently John Hawthorne (2004)—have 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 only use a proposition, *p*, as a premise in one's practical reasoning 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 supported by a relatively strong justification but which does not amount to knowledge.

In response to Hawthorne's claim about the pivotal role of knowledge in practical reasoning, Matt Weiner (2006) 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 on 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 on 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 the 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 is not vulnerable to criticism, a different result will be generated. (In the case under consideration, for example, it would be remiss to not check the lottery result in a local newspaper given that there is a chance that you’ve 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 than one knows one’s premises.

Whether or not Weiner is right about this, the possibility that he raises that the value that we place on knowledge might be due to the fact that it is a Swiss Army Concept in this way is intriguing, 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 beings with a thesis about the value of knowledge. In short, Craig hypothesises that knowledge is important to us because it fulfils the valuable function of enabling us to recognise 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.

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 more intellectually 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.¹⁵

7. CREDIT, LUCK, AND VALUE

We noted above that one of the key parts of the virtue-theoretic story of why knowledge is valuable is that knowledge is of credit to the agent. The guiding thought here is that knowledge is an achievement that is creditable to the agent, in the sense that the agent,

through her cognitive skills (i.e., her epistemic virtues, where these are construed broadly so that they could also include her cognitive faculties), gained a true belief *because* of her virtue. If knowledge is a distinctive achievement in this way—something that is worthy of credit—then this could, at least in part, account for the special value of knowledge.

This intuition that knowledge is always credit-worthy is often thought to be closely related to a second widely-held intuition about knowledge that it involves a non-accidental or non-lucky true belief. As the Gettier cases illustrate, one can't acquire knowledge by getting to the truth by luck. This sort of anti-luck intuition about knowledge might seem to suggest by itself that knowledge should be regarded as some kind of achievement, in the sense of being something that the agent can take credit for, since genuine achievements are precisely those sorts of things that are not gained by luck.

Interestingly, however, one can spell-out this anti-luck intuition in a way that does not necessarily implicate the 'knowledge as credit' thesis, but rather allows knowledge to be ascribed even in cases where the knowledge in question is not (or at least not obviously) of any credit to the agent. There are various ways of capturing this 'anti-luck' intuition, but most appeal to the so-called "safety" principle which, stated informally, demands that if one knows then one's true belief could not have easily been false.¹⁶ Various modal specifications of this principle have been offered, but we needn't get drawn into this issue here. What is important is rather that there are cases of true beliefs which could not have easily been false—on any intuitive cashing-out of that informal notion at any rate—but which do not seem to involve the agent accruing any credit for the true belief at issue. The question that faces us is thus whether we should count such cases as knowledge.

One such case is offered by Jennifer Lackey (*forthcoming*).¹⁷ 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 that she sees. Suppose the person that this agent 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. Nevertheless, as Lackey points out, it is far from obvious that it is of any credit to the agent in this case that she has formed a true belief (it might be of credit to the person she asks that she has a true belief about where Sears tower is, but that's another matter). More specifically, that the agent has a true belief in this case doesn't seem to reflect the operation of *her* epistemic virtues or cognitive faculties at all—which is what the main

proponents of the ‘credit’ thesis, virtue epistemologists, hold—but rather depends on the fact that the informant in question is suitably epistemically virtuous.

There are moves that those who hold to the credit thesis can make to this line of argument. For one thing, they can contend that in order for this to be a *bona fide* case of knowledge it is essential that the agent concerned is employing certain kinds of epistemic virtue. For example, presumably this agent wouldn’t just ask *any* person on the street—she wouldn’t ask someone who was wandering about aimlessly, incoherently muttering to herself, for example—and so this indicates a degree of sensitivity to the relevant factors in this regard.

The trouble with this sort of response is that it only shows, at best, that agents can sometimes know even while having a true belief that is of very minimal credit to her. The problem is, however, that if the credit involved is this minimal, then in what sense is the credit reflecting a cognitive achievement on the part of the agent? Indeed, one can easily multiply cases here. There can, for example, be instances of very basic beliefs formed via one’s perceptual faculties, even though one does not have any good reflectively accessible grounds to back-up those beliefs. If the beliefs are true and safe, then should they count as knowledge? It is hard to see why not if we allow credit to be a very minimal affair. After all, in some minimal sense, it is due to the cognitive traits of the agents that she is gaining a true belief, and this, according to the virtue-theoretic story, should suffice for knowledge.¹⁸ This suggests that the distinctive value of knowledge does not lie in the fact that knowledge is particularly credit-worthy, but rather relates to other features of knowledge, such as its immunity to epistemic luck.

In any case, recall the starting-point for our discussion about epistemic value—Plato’s reflections on the value of knowledge in the *Meno*. The main point that Plato raised in order to make sense of the greater value of knowledge over true belief was not to appeal to the credit that accrued to the knower, but rather concerned the ‘stability’ of the true belief itself; the fact that it was ‘tethered’ to the facts. This idea more naturally allies itself with the safety-based, and thus purely anti-luck, conception of knowledge than with the anti-credit conception of knowledge.

One natural way of reconciling the competing claims in this regard—put forward by Duncan Pritchard (*forthcomingb*)—would be to contend that what gives knowledge its

distinctive value over true belief relates to its immunity to luck, while also contending that if one's true belief is in addition of credit to one, then this raises the value of the knowing state (even though it is not essential to knowing that it be of credit to one, at least not to any substantive degree). In short, the hypothesis is that there are certain epistemic standings which are more robust than the epistemic standing required for merely knowing, where these elevated epistemic standings raise the value of the knowing state.

Independently of the plausibility of this idea, it is certainly reasonable to contend that there may be certain ways in which one can know that are more valuable than the kind of knowing that obtains where one merely meets the most minimal conditions. If understanding is a kind of knowing (or at least could be), then this could be one such elevated epistemic standing. This way of approaching the issue could thus engage with the related literature here on understanding that we surveyed above. The same goes for other elevated epistemic standings like wisdom.¹⁹ At the very least, one positive outcome of this sort of approach could be to shift the focus of contemporary epistemological theorising away from the merely minimal conditions for knowledge—a focus that has arisen largely in response to the Gettier problem—and move it towards higher epistemic standings.

8. SCEPTICISM AND EPISTEMIC VALUE

This issue of the differing value of certain kinds of knowing—along with the related issue that merely knowing might be valuable in quite a minimal fashion—could also have a bearing on the problem of scepticism. Oddly, there is, as yet, very little in the way of substantive discussion of the relationship between these two topics (and so this section will be brief).²⁰ The most obvious way in which these two topics relate is in terms of explaining the importance of scepticism. Most sceptical argument in the contemporary literature are directed at the possession of knowledge. If, however, knowledge turns out to be of little value to us, then why should we care whether or not scepticism about knowledge is true?

The two issues could also connect in more subtle ways, however. To begin with, certain kinds of responses to scepticism about knowledge proceed by, in effect, down-grading what is required of knowledge in order to evade the sceptical challenge. If knowledge is down-graded in this way, however, then it is natural to suppose that knowledge under this

construal is of less value. The question that emerges then, of course, is whether what we have rescued from the sceptic's grasp is really what is of value to us. Indeed, this way of approaching the issue can, I think, make it clear why certain commentators find one prominent style of response to the sceptical problem so unappealing.

Consider a formulation of the sceptical argument directed at knowledge which is common in the literature:²¹

- (1) I don't (indeed, am unable to) know the denials of sceptical hypotheses (e.g., that I am not a brain-in-a-vat).
- (2) If I don't know the denials of sceptical hypotheses, then I don't know any of the many 'mundane' propositions which I ordinarily take myself to know and which I know are incompatible with a sceptical hypothesis.
- (C) I don't know any of the many 'mundane' propositions which I ordinarily take myself to know.

A common move made by epistemic externalists in response to this sort of sceptical argument is to claim that provided one can have knowledge even while lacking good reflectively accessible grounds in favour of what one believes (as the externalist claims), then one can have knowledge of the denials of sceptical hypotheses, *contra* premise (1) in the argument just given. The rationale for this sort of move is that our intuition that we are unable to know the denials of sceptical hypotheses arises because we recognise that we lack good reflectively accessible grounds in favour of our beliefs in this respect. If we drop this requirement, says the externalist, then the way is opened for a very direct response to the sceptical problem.

Of course, the externalist still needs to tell a story about what it is virtue of that the agent knows in such cases, given that the story here does not advert to reflectively accessible grounds. One approach might be to say that the agent deduces this belief from her beliefs in the 'mundane' propositions which she knows are inconsistent with sceptical hypotheses, where these former beliefs are reliably formed. Thus, if one has a reliably formed true belief (and hence knowledge) that one is presently seated, for example, and one knows that this entails that one is not a brain-in-a-vat, then one can deduce that one is not a brain-in-a-vat, where the epistemic support that one has for the former belief transfers across to support the latter belief as well.

If one holds, in common with the theorists noted above, that knowing is always (to some substantive degree) of credit to one, then this avenue of response might not be all that

appealing. After all, it doesn't seem to be of any great credit to *you* that this belief in the inferred proposition is true, but rather seems to largely depend upon a kindness from the world (that you are, as a matter of fact, in an environment where the relevant belief-forming processes that you are employing are reliable). Put another way, the problem that arises here is that the 'knowledge' that is being rescued from the sceptic does not appear to be necessarily all that valuable.

Interestingly, our discussion in the last section on how Plato's response to the value problem in the *Meno* seems to demand something which is in principle weaker than the requirement that knowledge is always of credit to the agent concerned seems applicable here. In particular, a reliabilist account of knowledge—where the reliability condition is construed as one way of capturing the 'anti-luck' requirement on knowledge—could, as we have seen, allow knowledge in cases where the agent is not obviously deserving of any particular credit. As I also noted above, however, this thought is entirely consistent with the intuition that knowledge would be of more value to an agent if it is of more credit to that agent. One way of spelling out this idea thus treats knowledge as more valuable than mere true belief, but still not always quite as valuable as one might have otherwise thought.

Part of the attraction of construing the sceptical debate in these value-theoretic terms is that it might be able to cast light on a certain dissatisfaction that has been raised regarding externalist responses to scepticism. For example, this issue of epistemic value may well be part of what is underlying so-called 'metaepistemological' challenges to externalist responses to scepticism (e.g., Fumerton 1990; Stroud 1994), since such challenges seem to take the thought that 'knowledge' by externalist lights need not be a particularly valuable state as a starting-point from which to launch their critiques. More generally, there are good grounds for thinking that an investigation into how value claims are informing the sceptical debate could prove to be very enlightening.

9. THE VALUE OF TRUE BELIEF

So far, in common with most of the contemporary literature in this regard, we have tended to focussed on the value of knowledge relative to other epistemic standings. A related debate in this respect, however—one that has taken place often in tandem to the mainstream debate on

the value of knowledge—has specifically concerned itself with the value of true belief, and I want to close by making some relatively brief remarks about the state-of-play in this regard.

Few commentators treat truth or belief as being by themselves valuable,²² 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, is the use of the word ‘often’. After all, it is also often the case that a true belief might actually mitigate 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 more value—to have fewer true beliefs, and possibly more false ones, if this meant that the true beliefs that one had were regarding matters of real consequence.

At most, then, we can say that true beliefs often have instrumental value. What about intrinsic value? One might think that if the general instrumental value of true belief was moot then so would be the intuitively stronger thesis that true belief is generally intrinsically 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 in itself, regardless of what further prudential goals we might have.²³ 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* intrinsically good to believe the truth (cf. David 2003; 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* intrinsically good without being all things considered good. For example, it may be intrinsically 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 presently 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 intrinsically valuable. Why not instead just treat true belief as often of instrumental value and leave the matter at that? I think the answer to this 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 intrinsically valuable—and only typically instrumentally valuable—then this seems to down-play the status of the epistemological project.²⁴

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 intrinsically 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* intrinsically valuable. Sosa (see especially 2004*a*, but also 2001; 2003) seems to (almost) 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 intrinsically 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.²⁵ Kvanvig (see especially 2005) would be an example of someone who endorses epistemic value pluralism, since he thinks that there are number of fundamental epistemic goals, with each of them being of intrinsic 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 empirical adequate theory, making sense of the course of 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 only valuable because of their connection to the truth, and hence not fundamental epistemic goals at all.²⁶

Presumably, though, it ought to also be possible to make a case for an epistemic value pluralism where the fundamental epistemic goals were not intrinsically valuable (or, at least, *à la* Sosa, where one avoided taking a stance on this issue). More precisely, if an epistemic value monism which does not regard the fundamental epistemic goal as intrinsically 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 offered a plausible supporting story.²⁷

10. CONCLUDING REMARKS

I will close by outlining what I believe will be (or at least ought to be) the three main directions for future research in this area.

First, I think it is going to be increasingly important that those philosophers with a specific background in value theory are brought into this epistemological debate. When the discussion of epistemic value was in its infancy, it didn't matter quite so much that epistemologists were not as *au fait* with the nuanced material in value theory as they could be, since the points being made were of their nature quite general. Now that the argument has moved on to a more sophisticated level, however, it is vital that epistemologists are able to draw on the expertise of those with more familiarity with value theory.

Second, it is an oddity of the debate so far that much of the work on the value of truth (/true belief) has taken place independently from the debate regarding the value of other epistemic standings like knowing and understanding. More integration is required between these two realms of discussion, given their obvious relevance to each other.

Finally, as I argued in §8, I think there is much to be gained from reappraising the debate surrounding the sceptical problem in the light of a value-driven epistemology. In parts of this debate a kind of *impasse* has been reached, with both parties claiming that their opponents do not understand them and/or the issues. There is a real prospect here that re-

casting some of these disputes in value-theoretic terms could enable this *impasse* to be overcome.²⁸

REFERENCES

- Alston, W. (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation*, Cornell University Press, Ithaca, NJ.
- Axtell, G. (1997). 'Recent Work on Virtue Epistemology', *American Philosophical Quarterly* 34, 1-26.
- Brady, M. S. (2006). 'Appropriate Attitudes and the Value Problem', *American Philosophical Quarterly* 43.
- Brady, M. S., & Pritchard, D. H. (eds.) (2003). *Moral and Epistemic Virtues*, Blackwell, Oxford.
- Brentano, F. (1969). *The Origin of Our Knowledge of Right and Wrong*, (tr.) translated by R. Chisholm, Routledge & Kegan Paul, London.
- Brogaard, B. (2005). 'I Know. Therefore, I Understand', *typescript*.
- (Forthcoming). 'Can Virtue Reliabilism Explain the Value of Knowledge?', *Canadian Journal of Philosophy*.
- Brogaard, B., & Smith, B. (2005). 'On Luck, Responsibility and the Meaning of Life', *Philosophical Papers* 34.
- Chappell, T. (Forthcoming). 'Jonathan Kvanvig's *The Value of Knowledge and the Pursuit of Understanding*', *Religious Studies*.
- Chisholm, R. (1986). *Brentano and Intrinsic Value*, Cambridge University Press, Cambridge.
- Craig, E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Clarendon Press, Oxford.
- David, M. (2001). 'Truth as the Epistemic Goal', *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Virtue, and Responsibility*, (ed.) M. Steup, Oxford University Press, Oxford.
- (2005). 'Truth as the Primary Epistemic Goal: A Working Hypothesis', *Contemporary Debates in Epistemology*, (eds.) E. Sosa & M. Steup, 296-312, Blackwell, Oxford.
- DePaul, M. (1993). *Balance and Refinement: Beyond Coherence Methods of Moral Inquiry*, Routledge, New York.
- (2001). 'Value Monism in Epistemology', *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Virtue, and Responsibility*, (ed.) M. Steup, Oxford University Press, Oxford.
- DePaul, M., & Grimm, S. (Forthcoming). 'Jonathan Kvanvig's *The Value of Knowledge and the Pursuit of Understanding*', *Philosophy and Phenomenological Research*.
- Elgin, C. (Forthcoming). 'Is Understanding Factive?', *The Value of Knowledge*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, Oxford University Press, Oxford.
- Fumerton, R. (1990). 'Metaepistemology and Skepticism', *Doubting: Contemporary Perspectives on Skepticism*, (eds.) M. D. Roth & G. Ross, 57-68, Kluwer Academic Publishers, Dordrecht, Holland.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings*, Clarendon Press, Oxford.

- Goldman, A. (1999). *Knowledge in a Social World*, Oxford University Press, Oxford.
- 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 University Press, Cambridge..
- (2002). 'Knowledge as Credit for True Belief', *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) M. DePaul & L. Zagzebski, Oxford University Press, Oxford.
- (Forthcoming). 'The Value Problem', *The Value of Knowledge*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, Oxford University Press, Oxford.
- Grimm, S. (Forthcoming). 'Is Understanding a Species of Knowledge?', *British Journal for the Philosophy of Science*.
- Hawthorne, J. (2004). *Knowledge and Lotteries*, Oxford University Press, Oxford.
- 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. (1992). *The Intellectual Virtues and the Life of the Mind: On the Place of the Virtues in Contemporary Epistemology*, Rowman & Littlefield, Savage, Maryland.
- (1998). 'Why Should Inquiring Minds Want to Know?: Meno Problems and Epistemological Axiology', *The Monist* 81, 426-51.
- (2003). *The Value of Knowledge and the Pursuit of Understanding*, Cambridge University Press, Cambridge, UK.
- (2004). 'Nozickian Epistemology and the Value of Knowledge', *Philosophical Issues* 14, 201-18.
- (2005). 'Truth is not the Primary Epistemic Goal', *Contemporary Debates in Epistemology*, (eds.) E. Sosa & M. Steup, 285-96, Blackwell, Oxford.
- (Forthcoming). 'Responses', *The Value of Knowledge*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, Oxford University Press, Oxford.
- Lackey, J. (Forthcoming). 'Why We Don't Deserve Credit for Everything We Know', *Synthese*.
- Lynch, M. (2004). *True to Life: Why Truth Matters*, MIT Press, Cambridge, MA.
- (Forthcoming). 'The Values of Truth and the Truth of Values', *The Value of Knowledge*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, Oxford University Press, Oxford.
- McDowell, J. (1995). 'Knowledge and the Internal', *Philosophy and Phenomenological Research* 55, 877-93.
- Neta, R., & Pritchard, D. H. (Forthcoming). 'McDowell and the New Evil Genius', *Philosophy and Phenomenological Research*.
- Percival, P. (2003). 'The Pursuit of Epistemic Good', *Metaphilosophy* 34, 29-47; and reprinted in Brady & Pritchard (2003).
- Pritchard, D. H. (2002a). 'Resurrecting the Moorean Response to the Sceptic', *International Journal of Philosophical Studies* 10, 283-307.
- (2002b). 'Recent Work on Radical Skepticism', *American Philosophical Quarterly* 39, 215-57.
- (2003). 'Virtue Epistemology and Epistemic Luck', *Metaphilosophy* 34, 106-30; and reprinted in Brady & Pritchard (2003).
- (2005). *Epistemic Luck*, Oxford University Press, Oxford.

- (2006a). ‘Scepticism, Epistemic Luck, and Epistemic Value’, *typescript*.
- (2006b). ‘Virtue Epistemology and Epistemic Luck, Revisited’, paper presented at an invited symposium on *Intellectual Virtue* at the 2006 Pacific Division Meeting of the *American Philosophical Association*.
- (Forthcominga). ‘Anti-Luck Epistemology’, *Synthese*.
- (Forthcomingb). ‘Knowledge and Value’, *Think* 16; to be reprinted in *Epistemology (Royal Institute of Philosophy Lectures)*, (ed.) A. O’Hear, Cambridge University Press, Cambridge.
- Rabinowicz, W., & 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. (1998). ‘What are the “Chances” of Being Justified?’, *The Monist*, 81, 452-72.
- (2002). ‘Reliability and the Value of Knowledge’, *Philosophy and Phenomenological Research* 64, 79-96.
- (2003). ‘Understanding Virtue and the Virtue of Understanding’, *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) L. Zagzebski & M. DePaul, Oxford University Press, Oxford.
- (2005). ‘Review of *The Value of Knowledge and the Pursuit of Understanding*, by J. Kvanvig’, *Notre Dame Philosophical Reviews*, <http://ndpr.nd.edu/review.cfm?id=1393>.
- (2006). ‘The Value Turn in Epistemology’, *New Waves in Epistemology*, (eds.) V. Hendricks & D. H. Pritchard, Ashgate, Aldershot.
- (Forthcoming). ‘Getting the Meno Requirement Right’, *The Value of Knowledge*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, Oxford University Press, Oxford.
- Ryan, S. (1999). ‘What is Wisdom?’, *Philosophical Studies* 93, 119-39.
- Sainsbury, R. M. (1997). ‘Easy Possibilities’, *Philosophy and Phenomenological Research* 57, 907-19.
- Sartwell, C. (1992). ‘Why Knowledge is Merely True Belief’, *Journal of Philosophy*, 89, 167-80.
- Scanlon, T. M. (1998). *What We Owe to Each Other*, Harvard University Press, Cambridge, MA.
- Sosa, E. (1985). ‘Knowledge and Intellectual Virtue’, *The Monist* 68, 224-45.
- (1991). ‘Intellectual Virtue in Perspective’, *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge University Press, Cambridge.
- (2000). ‘How to Defeat Opposition to Moore’, *Philosophical Perspectives* 13, 141-54.
- (2001). ‘For the Love of Truth’, *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Virtue, and Responsibility*, (ed.) M. Steup, Oxford University Press, Oxford.
- (2003). ‘The Place of Truth in Epistemology’, *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) L. Zagzebski & M. DePaul, Oxford University Press, Oxford.
- (2004a). ‘Epistemic Normativity’, manuscript of paper delivered at the 2004 *Virtue Epistemology* conference, University of Stirling. (NB. An updated version of this paper appeared as part of Sosa’s John Locke Oxford lectures in 2005, and will be published with the rest of these lectures as part of a book).
- (2004b). ‘Reply to Linda Zagzebski’, *Sosa and His Critics*, (ed.) J. Greco, Blackwell,

- Oxford.
- Stroud, B. (1994). ‘Scepticism, ‘Externalism’, and the Goal of Epistemology’, *Proceedings of the Aristotelian Society* (supp. vol.) 68, 290-307.
- Swinburne, R. (1999). *Providence and the Problem of Evil*, Oxford University Press, Oxford.
- (2000). *Epistemic Justification*, Oxford University Press, Oxford.
- van Fraassen, B. (1980). *The Scientific Image*, Clarendon Press, Oxford.
- Weiner, M. (2006). ‘Does Knowledge Matter?’, *typescript*.
- Whitcomb, D. (2006). ‘Wisdom in Mind and Action’, *typescript*.
- Wiggins, D. (1987). ‘A Sensible Subjectivism?’, *Needs, Values, Truth*, Blackwell, Oxford.
- Williamson, T. (2000). *Knowledge and Its Limits*, Oxford University Press, Oxford.
- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge University Press, Cambridge, UK.
- (1999). ‘What is Knowledge?’, *Epistemology*, (eds.) J. Greco & E. Sosa, 92-116, Blackwell, Oxford.
- (2001). ‘Recovering Understanding’, *Knowledge, Truth, and Obligation: Essays on Epistemic Justification, Virtue, and Responsibility*, (ed.) M. Steup, Oxford University Press, Oxford.
- (2003a). ‘The Search for the Source of the Epistemic Good’, *Metaphilosophy* 34, 12-28; and reprinted in Brady & Pritchard (2003).
- (2003b). ‘Intellectual Motivation and the Good of Truth’, *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) M. DePaul & L. Zagzebski, 135-54, Oxford University Press, Oxford.
- (2004). ‘Epistemic Value Monism’, *Sosa and His Critics*, (ed.) J. Greco, Blackwell, Oxford.

NOTES

- ¹ For an excellent survey of recent work on virtue epistemology, see Axtell (1997).
- ² For three prominent examples of this sort of argument, see Zagzebski (1999), Williamson (2000, chapter 1), and Kvanvig (2003). I discuss this sort of argument again briefly below.
- ³ See also DePaul (1993), Jones (1997), Swinburne (1999; 2000), Riggs (2002), and Kvanvig (2003).
- ⁴ See especially Kvanvig (2003), who attributes the problem to Swinburne (1999; 2000).
- ⁵ Jones (1997) makes a similar point in his value-driven critique of reliabilism. In addition, he explicitly notes that other epistemological theories will be subject to the problem for reliabilism that he has identified, so long as, like reliabilism, they buy into this sort of epistemological picture, what he refers to as “epistemic instrumentalism”.
- ⁶ See also Brogaard & Smith (2005, §5).
- ⁷ A similar point is raised by Percival (2003) in his commentary on Zagzebski (2003a).
- ⁸ Brady attributes this thesis to Brentano (1969), and claims that it finds contemporary expression in the work of Chisholm (1986), Wiggins (1987), Gibbard (1990), and Scanlon (1998), amongst others.
- ⁹ For an excellent critical discussion of this assumption, see DePaul & Grimm (*forthcoming*).
- ¹⁰ In a similar fashion, there is a form of epistemic internalism put forward in the literature—by McDowell (e.g., 1992)—which offers an account of the internalist justification condition which is Gettier-proof, a possibility that Kvanvig does not consider. For more discussion of this aspect of McDowell’s position, see Neta & Pritchard (*forthcoming*).
- ¹¹ I’ve argued elsewhere—e.g., Pritchard (2003; 2005, chapter 7; 2006b)—that there is no virtue-theoretic response to the Gettier problem available. More recently, however, I’ve become less sure of this (though am still some way off being convinced that there is a response available).

¹² For an example of a philosopher of science who urges that we draw a distinction between acceptance and belief that functions in roughly this fashion, see van Fraassen (1980).

¹³ For more on Kvanvig's account of understanding, see DePaul & Grimm (*forthcoming*) in which they argue (§5), *inter alia*, that understanding can be Gettiered (contrary to what Kvanvig claims). See also Chappell (*forthcoming*) and Grimm (*forthcoming*). For more on the epistemology of understanding more generally, see Zagzebski (2001) and Riggs (2003).

¹⁴ A more extreme position is held by Sartwell (1992), who claims that knowledge is no more valuable than true belief. See also footnote 25 below.

¹⁵ See Jones (1997) for an interesting, and largely sympathetic, discussion of this aspect of Craig's project.

¹⁶ The literature on safety is itself now rather large. For some of the main texts advocating a version of the safety principle, see Sainsbury (1997), Sosa (1999), Williamson (2000), and Pritchard (2002a). See Pritchard (2005, part two; *forthcominga*) for a development of a safety-based account of knowledge which is explicitly developed in the light of the 'anti-luck' intuition just mentioned.

¹⁷ Note that Lackey, following the virtue epistemologists that she criticises, identifies the knowledge-entails-credit thesis with the knowledge-excludes-luck thesis. Accordingly, she holds that this sort of example in fact shows that we should abandon *both* theses. As I have noted (and will now show), however, we can separate out these intuitions and in doing so treat such cases as only raising problems for the first intuition.

¹⁸ Some virtue epistemologists get around this issue by insisting either that knowledge always requires a "perspective", so that it cannot be unreflective in this way—see, especially, Sosa (1991)—or that one can only gain knowledge through one's epistemic virtues, where these are now construed narrowly so that they don't include the cognitive faculties—see, especially, Zagzebski (1996).

¹⁹ There is a growing literature on wisdom, one that is starting to positively engage the related literature on epistemic value. See, for example, Ryan (1999) and Whitcomb (2006).

²⁰ I have begun to work on a paper on this topic myself—see Pritchard (2006a).

²¹ For a survey of recent work on radical scepticism, see Pritchard (2002b).

²² Though see Kvanvig (2003, chapter 1). Kvanvig's commitment to the value of truth and belief is criticized in DePaul & Grimm (*forthcoming*, §2).

²³ For examples of a train of argument of just this sort, see Goldman (1999, 3), Lynch (2004, 15-16), and Alston (2005, 31). I am grateful to Stephen Grimm for directing me to these references.

²⁴ Those sympathetic to this line of argument should, I think, be more tempted by the '*pro tanto*' conception of the intrinsic value of true belief than the '*prima facie*' conception. This is because on the former conception even when one's true belief is not all things considered valuable it is still the case that it has some degree of intrinsic value, while this would not be the case on the *prima facie* conception (since this value would be 'neutralised' by other considerations).

²⁵ Interestingly, Sartwell (1992) argues that since both knowledge and true belief are the goals of epistemic inquiry, it follows that they are identical. But this way of reasoning is only persuasive if one already has grounds to dismiss the possibility of epistemic value pluralism, and Sartwell doesn't himself offer such grounds.

²⁶ For further discussion of the prospects for epistemic value monism, see the exchange between Zagzebski (2004) and Sosa (2004b).

²⁷ Moreover, there may be a further axes along which this debate could be structured. For example, could one argue for a 'mixed' epistemic value pluralism, such that some of the fundamental epistemic goals are intrinsically valuable while some aren't? Such a view does not look *prima facie* very appealing, but no doubt conditions could be offered in its favour.

²⁸ For valuable discussions, and in many cases for allowing me to take a look at unpublished work, I am grateful to Jason Baehr, Michael Brady, Berit Brogaard, Tim Chappell, Michael DePaul, Catherine Elgin, John Greco, Stephen Grimm, Adrian Haddock, Chris Hookway, Carrie Jenkins, Ward Jones, Jon Kvanvig, Jennifer Lackey, Michael Lynch, Alan Millar, Daniel Nolan, Christian Pillar, Wayne Riggs, Robert Roberts, Barry Smith, Ernie Sosa, Finn Spicer, Matt Weiner, Dennis Whitcomb, Jay Wood, René van Woudenberg, and Crispin Wright.