THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Knowledge, Understanding and Epistemic Value

Citation for published version:

Pritchard, D 2009, 'Knowledge, Understanding and Epistemic Value', *Royal Institute of Philosophy Supplements*, vol. 64, pp. 19-43. <https://doi.org/10.1017/S1358246109000046>

Digital Object Identifier (DOI):

[10.1017/S1358246109000046](https://doi.org/10.1017/S1358246109000046)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Royal Institute of Philosophy Supplements

Publisher Rights Statement:

© Cambridge University Press. Pritchard, D. (2009). Knowledge, Understanding and Epistemic Value. Royal Institute of Philosophy Supplement, 64, 19-43doi: 10.1017/S1358246109000046

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Royal Institute of Philosophy Supplement

<http://journals.cambridge.org/PHS>

Additional services for ***Royal Institute of
Philosophy Supplement***:

Email alerts: [Click here](#)

Subscriptions: [Click here](#)

Commercial reprints: [Click here](#)

Terms of use : [Click here](#)

Knowledge, Understanding and Epistemic Value

Duncan Pritchard

Royal Institute of Philosophy Supplement / Volume 64 / July 2009, pp 19 - 43

DOI: 10.1017/S1358246109000046, Published online: 27 May 2009

Link to this article: http://journals.cambridge.org/abstract_S1358246109000046

How to cite this article:

Duncan Pritchard (2009). Knowledge, Understanding and Epistemic Value. Royal Institute of Philosophy Supplement, 64, pp 19-43 doi:10.1017/S1358246109000046

Request Permissions : [Click here](#)

Knowledge, Understanding and Epistemic Value

DUNCAN PRITCHARD

1.

It is a widespread pre-theoretical intuition that knowledge is distinctively valuable. If this were not so, then it would be simply *mysterious* why knowledge has been the focus of so much of epistemological theorising, rather than some other epistemic standing like justified true belief. Given this fact, however, it is obviously important to a theory of knowledge that it is able to offer a good explanation of why we have this intuition. Indeed, some, such as Jonathan Kvanvig (2003) and Timothy Williamson (2000), have argued that if a theory of knowledge does not make it transparent why knowledge is distinctively valuable then this is a decisive strike against it. We do not need to go this far, however. What is important is just that a theory of knowledge is able to adequately account for this intuition.

One very direct way of accounting for the intuition would be to offer a theory of knowledge which demonstrated why knowledge is distinctively valuable in the manner that we intuitively suppose. We will call proposals of this sort *validatory*, since they aim to validate our pre-theoretical intuitions about the value of knowledge. Positions of this sort have been offered by, for example, Linda Zagzebski (1996; 1999; 2003) and John Greco (2002; 2007; *forthcominga*), and we will consider one such proposal in this respect below.

Notice, however, that one does not need to validate an intuition in order to account for it. One could instead put forward a theory of knowledge on which knowledge is not distinctively valuable, but which could explain why we might pre-theoretically *think* that knowledge is distinctively valuable. We will call proposals of this sort *revisionist*, since they revise our pre-theoretical intuitions about the value of knowledge. Mark Kaplan (1985), for example, famously argued that the moral of the post-Gettier literature was that it is not knowledge which is distinctively valuable but rather justified true belief – knowledge being justified true belief plus an anti-Gettier condition – but that since justified true belief usually sufficed for knowledge, the mistake was entirely natural. A second proposal

along these lines, which we will look at in more detail below, is offered by Kvanvig (2003) who argues that it is understanding, not knowledge, which is distinctively valuable, where understanding is an epistemic standing that is closely related to knowledge.

Of course, a final option in this regard is to simply argue that our intuitions on this score are simply wrong on closer analysis. That is, that there is no distinctively valuable epistemic standing. We will call proposals of this sort *fatalist*, since they do not hold out any hope of doing justice to our pre-theoretic intuitions about the value of knowledge in the way that revisionist proposals do. If you think, like Crispin Sartwell (1992), that knowledge just is true belief then you will probably be sympathetic to a view of this sort.¹

Clearly, a fatalist proposal will be by its nature an uncomfortable position to defend. In this paper I will be exploring a version of validationism and a version of revisionism, and along the way trying to avoid fatalism. As we will see, although both of the proposals that we will be looking at are problematic, a third position will emerge from our discussion which can at least offer us a plausible revisionist account.²

2.

The first response that we will be looking at is a form of validationism and it arises out of a certain virtue-theoretic account of knowledge. *Modest* virtue epistemological theories – of the sort defended by, for example, Greco (1999; 2000) in his early work – demand that a necessary condition of knowledge is that the agent forms her true belief via the stable and reliable cognitive abilities that make up her cognitive character. There is obviously a lot to be said about how a proposal of this sort is to be construed. One might build quite a lot into the notion of a cognitive ability, for example, or into the notion of a cognitive character. Depending on how one developed these notions the view could thus be more or less restrictive as an account of knowledge. We can set these issues to one side, however, since, as we will see, what is important for our purposes is the *structure* of a proposal of this sort. The rationale for adding this requirement to a theory of knowledge is that what we primarily want from

¹ S. Stich 1990 is also often credited with advancing a form of fatalism of this sort about the value of knowledge.

² For more on the problem of accounting for epistemic value, see Pritchard (2007b; cf. Pritchard 2007d).

Knowledge, Understanding and Epistemic Value

such a theory is an account of how we are being suitably sensitive to the facts when we know, and this makes cognitive abilities central to knowledge.

To see this point, consider the following example. Imagine that someone is in a room and forming her beliefs about the temperature of the room by looking at a thermometer on the wall. Suppose further that this is indeed a highly reliable way of forming beliefs in this regard, in the sense that every time she forms her belief in this way the belief so formed is true. Here's the twist. The thermometer is, unbeknownst to the agent, broken, and is fluctuating within a given range. This does not undermine the reliability of the agent's beliefs, however, for the simple reason that there is someone hidden in the room who is altering the thermostat in such a way as to ensure that every time the agent forms a belief about the temperature of the room by looking at the thermostat, her belief is true.

The agent in this case clearly does not have knowledge. Moreover, the right diagnosis of why the agent doesn't know is that the reliability of her belief-forming processes does not reflect any cognitive ability on her part. It is not as if she is being sensitive to the facts in the way that she is forming her beliefs, but rather that the facts are being sensitive to her beliefs – i.e. the direction of fit is all wrong. Virtue epistemology offers a straightforward way of dealing with cases like this, since the fact that the agent is not forming her true belief via her cognitive abilities suffices to entail, on this view, that she lacks knowledge.

Crucially, though, standard forms of virtue epistemology do not regard this appeal to cognitive ability as sufficing to offer a complete account of knowledge. This is because of Gettier-style cases, cases in which something intervenes 'betwixt belief and fact'.³ Suppose that our agent is looking into a field and, using her reliable cognitive abilities, forms the belief that there is a sheep in the field. Suppose further that this belief is true, but that the agent is not in fact looking at a sheep but a big hairy dog which looks just like a sheep, and which is obscuring from view the sheep that is in the field. The agent in this case clearly lacks knowledge since it is just a matter of luck that her belief is true.⁴ Nevertheless, she is forming a true belief via the stable and reliable cognitive abilities that make up her cognitive character.

The standard way in which virtue epistemologists deal with Gettier-style cases is by supplementing the view with an anti-luck

³ I owe this way of putting the kind of luck in play in Gettier-style cases to P. Unger (1968).

⁴ This example is due to R. Chisholm 1977: 105.

condition, like the safety principle. This is the move made until quite recently by Greco (1999; 2000), for example (though, as we will see in a moment, Greco takes a very different line on this issue now).⁵ We do not need to get into the details of what is involved in a principle like safety here; what is important is just that such a condition ensures that the agent could not have easily been wrong, thereby denying knowledge to agents in the Gettier-style cases.

3.

More recent virtue-epistemic proposals have not taken this line, however, and have instead followed Ernest Sosa (1988; 1991; 2007) in arguing that if we ‘beef-up’ the ability condition on knowledge then we can deal with Gettier-style cases without appeal to an anti-luck condition. Significantly for our purposes, such a robust virtue epistemology is also able, or so the argument goes at any rate, to account for the distinctive value of knowledge.

Proponents of a robust virtue epistemology of this sort – such as Zagzebski (e.g. 1995) and, in more recent work, Greco (2002; 2007; *forthcominga*; *forthcomingb*) – argue that where modest virtue epistemology goes wrong is by simply requiring the conjunction of cognitive ability and cognitive success (i.e. true belief). So construed, it is possible for something to come ‘betwixt’ the cognitive ability and the cognitive success such that the success is ‘gettierized’. However, we can avoid this situation, they argue, so long as we require not just the conjunction of cognitive success and cognitive ability, but in addition demand that the cognitive success be *because of* cognitive ability, in the sense that the cognitive ability best explains the cognitive success.

Consider again the ‘sheep’ Gettier-style case described above and suppose that we add the ‘because of’ requirement. This certainly does seem to deal with this example since while the agent’s cognitive success arises out of her cognitive ability, the cognitive success is not *because of* her cognitive ability but rather because of some incidental fact about the environment (i.e. that there happened to be a sheep hidden from view behind the big hairy dog).

So it does seem as if this proposal can indeed deal with Gettier-style cases, and if it can then there is no need to add an anti-luck condition. Of course, adding this requirement will create

⁵ For more on safety, see Sosa 1999 and Pritchard 2002; 2005: ch. 6; 2006; 2007a.

problems in other respects, largely due to the fact that it is a complex matter offering the right account of the ‘because of’ relation. Indeed, on the standard view, the right semantics for causal explanation sentences is along contextualist lines, and this would seem to suggest that a robust virtue epistemology should be allied to a form of attributer contextualism. Although Greco (e.g. *forthcomingb*), for one, has embraced this consequence of his view, this is certainly a surprising alliance. Still, we needn’t get bogged-down in this issue.⁶ Let us take it that we have at least an intuitive sense of how to read these ‘because of’ claims. As we will see, the issues that we need to consider in this regard trade on examples where our intuitions are pretty clear-cut, and thus we ought to be able to ignore these complications without too much concern.

4.

The manner in which a proposal of this sort can enable us to deal with the value problem is because knowledge on this view can plausibly be regarded as a type of achievement, and achievements in turn are often thought to be distinctively valuable. Let us consider the notion of an achievement first. Proponents of robust virtue epistemology maintain that an achievement is a success that is because of ability. Since knowledge on their view is to be understood as a cognitive success that is because of cognitive ability, that makes knowledge a cognitive achievement.

In order to see that this is a plausible account of achievement, imagine someone with a bow and arrow selecting a target, firing at that target, and hitting the target. Suppose, however, that the agent in question did not have any ability in this regard. Clearly, in such a case we would not credit the agent with an *achievement*, since it was just dumb luck that she was successful. *A fortiori*, if there is no ability involved then it cannot be because of such ability that the agent is successful and so the account of achievement on offer deals with such cases. Now suppose that we have an agent selecting a target and *skillfully* hitting that target with her arrow. Imagine, however, that the agent’s success is gettierized, in that something intervenes ‘betwixt’ ability and success. Perhaps, for example, a freak gust of wind blows the arrow off course, and then a second gust of wind blows it back on course again. By anyone’s lights,

⁶ I discuss Greco’s view in this regard in more detail in Pritchard (*forthcomingb*).

although the relevant ability is present, we would not say that this success is because of the agent's ability since it is clearly due to the fortuitous second gust of wind. Moreover, by the same token, we would not regard this as an achievement on the part of the agent either, since the success is not properly creditable to her. Again, the account of achievement under consideration deals with such cases. Finally, suppose that the agent's success with the arrow arises out of the relevant skill and is not gettierized. Surely we would now say that this success is because of the archer's ability and, crucially, we would also treat such a success as an achievement, just as the account of achievement on offer would predict.

So there is a good case to be made for thinking that achievements are successes that are because of ability, and if this claim is allowed then the thesis that knowledge is a type of achievement – a cognitive achievement – follows immediately on the robust virtue epistemic view. With this in mind, let us now examine the further claim that achievements are distinctively – indeed, *finally* (i.e. non-instrumentally) – valuable. To begin with, notice that from a practical point of view it might not matter whether or not a success is because of ability, and so constitutes an achievement. If hitting that target wins you the competition, for example, then it may not matter to you whether the success in question was, say, gettierized. Nevertheless, we do value achievements very differently from successes that fall short of being achievements, as when they are gettierized or are due to dumb luck rather than ability. In particular, a genuine achievement seems to be valuable in its own right, independently of any practical value the success in question might generate. For example, all other things being equal, we would surely think that it is better to hit the target because of one's skill than not, even if there is no instrumental value from exhibiting an achievement in this case. This seems to suggest that achievements are *finally* valuable.

If this is right, and knowledge is a type of achievement, then it seems that it will inherit the value of achievements. The reason why knowledge is distinctively valuable is because knowledge is an achievement and achievements are distinctively valuable. This would be a very neat response to the value problem. Moreover, notice that this would be a case in which one's theory of knowledge makes it explicit just why we care about knowledge in the way that we do. As a validationist response to the problem of epistemic value, it is thus very attractive.⁷

⁷ Interestingly, proponents of this thesis always express the view as being that knowledge is *intrinsically* valuable. Given that the

5.

Now there clearly are some *prima facie* objections to the idea that achievements are finally valuable. Some achievements, after all – such as easy, trivial or wicked achievements – do not seem to be very valuable at all. Notice, however, that the claim is only that achievements are finally valuable *qua* achievements; the thesis is not that all achievements are of *overall* – i.e. all things considered – value. It is thus entirely open to the defender of this thesis to maintain that the overall value of lots of achievements is very low – perhaps negative, if you believe that such a thing is possible – even while defending the specific thesis that achievements are finally valuable.

Alternatively, the proponent of such a thesis could argue for a modified version of the thesis along more holistic lines by saying, for example, that it is in the *nature* of achievements to be finally valuable, even though some achievements, because of their other properties, are not finally valuable.⁸ For our purposes, so long as cognitive achievements are the kind of thing that is finally valuable, then that would probably suffice to ensure the distinctive value of knowledge. The thesis that knowledge is distinctively valuable surely does not require us to claim that *all* knowledge is distinctively valuable. So long as it is in the nature of knowledge, *qua* cognitive achievement, to be finally valuable, then that would almost certainly suffice.

In any case, we can set these issues to one side just now, since there are surely good *prima facie* grounds for thinking that the claim that achievements are distinctively valuable can be adequately motivated. Moreover, the objection I want to raise to this validationist account of epistemic value does not turn on any qualms about the value of achievements but rather concerns the thesis that knowledge should be understood as a cognitive achievement. As we will see, there are good reasons for thinking that such a thesis is unsustainable.

non-instrumental value of the cognitive success is due to the relational properties of that success – i.e. due to how that success was achieved – it should be clear that it is specifically final value that is in play here. For more on the distinction between intrinsic and final value, see Rabinowicz & Roennow-Rasmussen 1999; 2003. For more on the distinction as it applies to the debate regarding epistemic value, see Pritchard 2007*d*, B2.

⁸ Thanks to Mike Ridge for this suggestion.

6.

Think again about the case of the archer described above. As before, suppose that this archer selects a target at random and skilfully fires at that target, hitting the target as a result of her skill – that is, nothing intervenes ‘betwixt’ ability and success such that the success is gettierized. We noted above that in such a case we would surely regard the agent’s success as being because of her ability, and therefore credit her with a genuine achievement.

Suppose that we add a further twist to this case, however, and stipulate that had the agent chosen any other target on that range then she would have missed because, unbeknownst to her, the targets in question have a forcefield in them which would have deflected the arrows. Her success is thus lucky in the sense that she could have very easily been unsuccessful. Nevertheless, does luck of this sort undermine the agent’s achievement? I say not. Indeed, achievements seem entirely compatible with luck of this sort, unlike the Gettier-style luck which intervenes between ability and success. After all, the agent really is hitting the target because of her ability, and the luck in question – which we might term ‘environmental’ luck – does nothing to undermine this.

Insofar as we grant this point, however, then it creates problems for the knowledge-as-achievement thesis. After all, we can construct an example which is structurally analogous to the one just given but where the environmental luck in question does undermine the agent’s putative knowledge. The famous ‘barn facade’ example is the best illustration of this point. Here we have an agent who sees a barn in clear daylight and so forth and, using her reliable cognitive abilities, forms a belief that what she sees is a barn. Moreover, this belief is true and is not gettierized since she really is looking at a barn (and thus nothing intervenes ‘betwixt’ belief and fact). Nevertheless, her true belief is epistemically lucky – in the sense that she could have easily been wrong – because unbeknownst to her she is in barn facade county where nearly all the barn-shaped objects are in fact fake barns which are indistinguishable to the naked eye from the real thing. Does our agent know that what she sees is a barn? Surely not, since her true belief is epistemically lucky – she could very easily have been wrong. But is her true belief, her cognitive success, a cognitive achievement? Well, if the ‘archer’ case just described is anything to go by, then it surely is. After all, her true belief really is because of her cognitive ability, so if that is what constitutes an achievement – and we have seen that there is good reason to think that this is the right way to think

Knowledge, Understanding and Epistemic Value

about achievements – then we should regard the agent’s cognitive success in this case as a cognitive achievement.

It thus seems that sometimes at least there is more to knowledge than a mere cognitive achievement, and this means that the knowledge-as-achievement thesis is false. In particular, it seems that exhibiting a cognitive achievement does not suffice to eliminate all the knowledge-undermining kinds of epistemic luck, such as an environmental epistemic luck.⁹

As far as I am aware, Greco is the only one to have engaged with an objection of this sort – in print, at any rate – and at different points he has made different responses to this objection to his view.¹⁰ Initially, his line was to deny that the agent in the ‘barn facade’ case – and thus, presumably, the agent in the ‘archer’ case also – had the ability in question.¹¹ Abilities are, after all, relative to environments, so this line of argument is not entirely outlandish. Nevertheless, we surely do not want the relativisation of abilities to environments to have the result that abilities must be infallible, and this seems to be the consequence of taking this line (for if the ability is not infallible then there is bound to be a case that we can construct in which the environment is such that the agent could very easily have been wrong). Moreover, neither do we want abilities to be construed in an unduly fine-grained manner such that the relevant ability is lost as soon as one enters the ‘deceptive’ environment in question.

A second line that Greco (*forthcomingb*) pushed was simply to insist that to say that a success was because of ability is *thereby* to say that it is not due to luck. This is not a helpful suggestion, however, since, as we have seen, while it is plausible to think that the ‘because of’ eliminates the kind of Gettier-style luck that intervenes ‘betwixt’ success and ability, it is actually far from obvious – and, indeed, counterintuitive – that it eliminates the very different sort of environmental luck at issue in the cases just given. It seems, then, that one cannot evade a problem like this through stipulation.

Ultimately, however, the line that Greco (2007b) has taken to this problem has been to offer further theses regarding the function of our concept of knowledge in order to explain why knowledge, *qua* cognitive achievement, should be more resistant to luck than other

⁹ For more on this point as it applies to Sosa’s view in particular, see Pritchard 2007c.

¹⁰ Note that I am not suggesting here that these responses are necessarily in tension with one another; indeed, there is every reason to think that they are complementary.

¹¹ See Greco *forthcominga*: §5.

types of achievement. This move will deal with the problem, but it does beg the question of whether a better way to deal with these cases would be to abandon the knowledge-as-achievement thesis rather than make exceptions for the case of knowledge that do not apply elsewhere.¹²

7.

This problem is even more pressing once one notices that there are cases of knowledge which, intuitively, do not involve cognitive achievements, so that the separation between knowledge and cognitive achievement goes in both directions. Consider the following case, originally offered by Jennifer Lackey (2007), albeit to illustrate a slightly different point.

Imagine our agent getting off the train in an unfamiliar town and asking the first person that she sees for directions. Suppose further that the informant does indeed have first-hand knowledge in this regard and communicates this information to our agent who subsequently heads off to her destination. We would naturally describe such a case as one in which the informant's knowledge was straightforwardly communicated to our agent; indeed, if we don't allow knowledge in cases like this then it seems that quite a lot of our putative testimony-based knowledge is called into question.

Crucially, however, it does not seem at all right to say that our agent's cognitive success is because of *her* cognitive ability. Indeed, the right thing to say seems to be that it is because of the *informant's* cognitive ability, or at least because of their combined cognitive efforts. But that means that sometimes knowledge requires a lot less than a cognitive achievement, contrary to the knowledge-as-achievement thesis.

It is important to be clear what the target of this objection is. Lackey (2007) herself takes it to show that one can have knowledge without it being of *any* credit to one that one has a true belief. But examples like this surely do not license this rather radical conclusion. After all, the agent in this case is exhibiting quite a lot of cognitive ability if one examines the case a little more closely. Although she asks the first person she meets, she wouldn't have asked just *anyone* (or *anything*). She wouldn't have asked a child, for example, or someone who was clearly a tourist (and she certainly wouldn't have asked a lamppost or a passing dog). Moreover, she is presumably

¹² I discuss Greco's treatment of these issues further in Pritchard *forthcomingb*. See also Kvanvig *forthcoming*.

sensitive to potential defeaters. If the informant had given her directions which were obviously fake, for example, then we would have expected her to have spotted this. Indeed, it is only if the agent is exercising her cognitive abilities in this way that it seems permissible to credit her with knowledge.

Nevertheless, the point remains that it is not because of her cognitive abilities that she is cognitively successful – by anyone's lights – even though it is of some credit to her that she is cognitively successful. But this is the point that we need to undermine, the knowledge-as-achievement thesis, since it demonstrates that there is sometimes less to knowledge than a cognitive achievement.

8.

If this were the only problem facing the knowledge-as-achievement thesis, then one might reasonably take the heroic route of denying the intuition in this case and insisting that it is because of the agent's cognitive abilities that she is cognitively successful. But once one combines this objection with that noted earlier – which demonstrates that there is sometimes *more* to knowledge than a cognitive achievement – then this points towards a different way of understanding knowledge.

Indeed, I would argue that what cases like this show is that the modest virtue epistemic proposal is preferable to the robust virtue epistemic proposal. After all, on this view the ability condition on knowledge is not 'beefed up' to the extent that knowledge demands a cognitive achievement, and so one does not get the problem posed by Lackey-style cases. Moreover, since there is also the anti-luck condition on knowledge, expressed in terms of a safety principle, then cases like the barn facade case, in which there is a cognitive achievement but also knowledge-undermining epistemic luck, are also dealt with.

Furthermore, notice that such a view is not necessarily in conflict with the story told by robust virtue epistemologists regarding the distinctive value of knowledge. At the very least, the modest virtue epistemic proposal is consistent with a revisionist response to the problem of epistemic value which says that it is not knowledge, strictly speaking, which is distinctively valuable, but rather cognitive achievement, an epistemic standing which (it seems) only comes apart from knowledge in peripheral cases.

But the modest virtue epistemic proposal might also be compatible with a validationist response to the problem of epistemic value as well.

If it is indeed true that knowledge and cognitive achievement only come apart in peripheral cases – and whether one finds this claim plausible may depend, in part, on one's wider epistemological theory – then one could argue that it is of the nature of knowledge to be distinctively valuable, even though it isn't always distinctively valuable. Perhaps, for example, all paradigm cases of knowledge are also cases of cognitive achievement.¹³ If that's right, then there might be scope to argue that knowledge is distinctively valuable after all.

9.

Let us put this tentative conclusion to one side for now, however, because I want to consider a second account of epistemic value which, as we will see, is relevant in this regard. This is the proposal that *understanding* is a distinctively valuable epistemic standing, a thesis which is often supplemented with the further claim that it is the *only* distinctively valuable epistemic standing, thereby making the view a form of revisionism.

It is easy to see the attraction of such a view, in that understanding does seem to be particularly valuable to us. More specifically, insofar as knowledge and understanding do indeed come apart, then understanding seems to be preferable to knowledge. As we might be tempted to put the point, we would surely rather understand than merely know.

Before we can evaluate a claim of this sort, however, we need to be a little clearer about what we are talking about. One problem that afflicts any direct comparison between knowledge and understanding is that knowledge (of the propositional sort that we are concerned with at any rate) is concerned with *propositions*, whereas understanding usually isn't, at least not directly. Interestingly, where understanding is of a proposition, it does seem to be pretty much synonymous with knows. On discovering that my train has been cancelled, I may well say to the person at the ticket office that I understand that the train is cancelled in such a way that I could just as well have used 'know' without any loss. If anything, using 'understand' in this way seems to weaken the effect of the assertion. If I say to you that I understand that you are angry with me then this has the positive effect of being a little less confrontational than a straight assertion that I know that you are angry with me (for one thing, it gives you the option to deny this without obviously accusing me of any ignorance).

¹³ I am grateful to Chris Hookway for this suggestion.

Knowledge, Understanding and Epistemic Value

Most uses of ‘understands’ are not like this, however. I want to take the paradigm usage of ‘understands’ to be in a statement like ‘I understand why such-and-such is the case’. Notice that this usage is very different from a more holistic usage which applies to subject matters, as in ‘I understand quantum physics’, or even ‘I understand my wife’. I think the holistic usage of ‘understands’ is related to the non-holistic, or atomistic, usage that is our focus, but the former raises problems of its own that we’ve not the space to cover here (though we will flag some of these problems as we go along).¹⁴

Regarding understanding-why – henceforth just ‘understanding’ – there are, interestingly, two standard views – a standard view within epistemology and a standard view *outside* of epistemology (particularly in the philosophy of science). The standard view within epistemology is that understanding is distinctively valuable but that it is *not* a species of knowledge. One finds a view of this sort in the work of such figures as Kvanvig (2003), Zagzebski (2001), Wayne Riggs (*forthcoming*) and Catherine Elgin (1996; 2004; *forthcoming*), and we will examine the motivation for such a thesis in a moment.

In contrast, outside of epistemology the consensus is clearly that understanding *is* a species of knowledge. In particular, most philosophers of science who have expressed an opinion on this matter have endorsed the claim that understanding why such-and-such is the case is equivalent to knowing why such-and-such is the case, where this is in turn equivalent to knowing that such-and-such is the case because of such-and-such. So, for example, my understanding of why my house burned down is equivalent to my knowing why my house burned down, where this in turn is tantamount to my knowing that my house burned down because (say) of faulty wiring. One finds a view of this sort – expressed in varying levels of explicitness – in the work of such figures as Peter Achinstein (1983), Wesley Salmon (1989), Philip Kitcher (2002), James Woodward (2003) and Peter Lipton (2004).¹⁵

¹⁴ For more on holistic and non-holistic conceptions of understanding, see B. Brogaard 2007.

¹⁵ Consider the following remark made by Lipton 2004: 30 and quoted in S. Grimm 2006: 1, for example: “Understanding is not some sort of super-knowledge, but simply more knowledge: knowledge of causes”. The natural way to read this passage is as suggesting that understanding why one’s house burned down is just knowing why it burned down – i.e. knowing that it burned down because of (say) faulty wiring. I am grateful to Grimm (2006) for alerting me to some of these references.

As we will see, I want to claim that both of these conceptions of understanding are wrong, at least strictly speaking, and that once we get clear on the relationship between understanding and knowledge we can make some progress towards dealing with the problem of epistemic value.

10.

Let us look first at some the accounts of understanding offered by epistemologists. One guiding theme in this discussion is that understanding is construed along epistemically internalist lines. One extreme example of this can be found in the work of Zagzebski (2001). She argues, amongst other things, that understanding is, unlike knowledge, “transparent” in the sense that there is no gap between seeming to understand and understanding. Relatedly, she also claims that understanding is, unlike knowledge, non-factive, in that even if one’s relevant beliefs were false, one’s understanding could be unaffected.¹⁶ Finally, she holds that understanding, unlike knowledge, is immune to epistemic luck, in that if one’s understanding is subject to such luck it will not thereby be undermined.

Of these claims, the first is clearly the most radical and also, I venture, the one that is most obviously false. To construe understanding in this way seems to reduce it to nothing more than some sort of minimal consistency in one’s beliefs, something which might well be transparent to one (though I’m actually doubtful of this). Understanding clearly involves much more than this, however. To see this, let us focus on the non-factivity claim that Zagzebski makes. This claim is also, I will argue, false, but if understanding does imply factivity in the relevant sense, then it will be easy to show that understanding is not transparent in the way that Zagzebski suggests.

To illustrate this point, consider my understanding of why my house has burned down. Let us grant the plausible assumption that this understanding involves a coherent set of relevant beliefs concerning, for example, the faulty wiring in my house. But now suppose that these beliefs are mistaken and that, in particular, there was no faulty

¹⁶ Riggs (*forthcoming*) and Elgin (*forthcoming*; cf. Elgin 1996; 2004) also argue that understanding is not factive, although their claim is ultimately much weaker than Zagzebski’s since it in effect only applies to certain conceptions of understanding (and not, in particular, to the non-holistic conception of understanding in play here).

wiring in my house and so it played no part in the fire. Would we still say that I understand why my house burned down? I think not. For sure, I *thought* I understood – indeed, it could well be that I *reasonably* thought that I understood – but the fact remains that I did not understand. Once one grants that understanding is factive in this way, however, then the transparency claim starts to look equally suspect, since if understanding is factive then it clearly cannot be transparent as the factivity of understanding would require there to be a distinction between thinking that certain facts obtain and their obtaining, contrary to what the transparency thesis demands.

So the transparency and non-factivity claims that Zagzebski offers are false. It is difficult to diagnose why Zagzebski made this mistake. Part of the reason may be that there is a failure to be clear about the type of understanding under consideration. After all, when it comes to the kind of holistic understanding that applies to a subject matter, this plausibly *is* compatible with at least *some* false beliefs about that subject matter, but this sort of understanding is precisely not the sort at issue. Moreover, it would seem that the analogue of Zagzebski's non-factivity claim as regards understanding when it comes to holistic understanding would be that such understanding can be possessed even though one has *no* relevant true beliefs, and that is surely implausible.¹⁷

More generally, however, I think the diagnosis for where Zagzebski's conception of understanding goes awry lies in overstating the internalist aspect of understanding. Understanding clearly is very amenable to an account along internalist lines, in the sense that it is hard to make sense of how an agent could possess understanding and yet lack good reflectively accessible grounds in support of that understanding. Understanding thus cannot be 'opaque' to the subject in the way that knowledge, by externalist lights at least, can sometimes be. Granting this, however, does not entail that one should regard understanding as non-factive, much less transparent.

¹⁷ It should be noted that there are some good arguments offered by Elgin (*forthcoming*) in this respect regarding the growth of understanding within false scientific theories, and the use of idealisations in scientific thinking, which might seem to suggest a conception of holistic understanding which is entirely non-factive. It would take us too far afield to consider these arguments, however, and Zagzebski clearly doesn't have considerations like this in mind when she offers her conception of (non-holistic) understanding. For my own part, I think that even here we should say that genuine understanding entails a system of beliefs which is broadly correct, at least as regards the beliefs that are fundamental to that system. For more on this point, see Pritchard 2007b: B5.

11.

With this in mind, let us consider a second account of understanding in the epistemological literature – due to Kvanvig (2003) – which does not succumb to the mistakes made by Zagzebski's account. Zagzebski holds that both knowledge and understanding are distinctively valuable. In contrast, Kvanvig maintains that it is only understanding that is distinctively valuable, where understanding is distinct from knowing.

Unlike Zagzebski, Kvanvig does not hold that understanding is transparent or non-factive. He does, however, treat the notion along internalist lines which, as we've just noted, is entirely proper. The way in which he distinguishes knowledge from understanding is primarily through two further claims. The first is that understanding, unlike knowledge, admits of degrees. The second is that understanding, unlike knowledge, is immune to epistemic luck, a thesis which we saw Zagzebski putting forward a moment ago.

The import of the first claim is, I think, moot. After all, even if this is true, it needn't follow that there are cases of knowledge which aren't corresponding cases of understanding, or that there are cases of understanding which aren't corresponding cases of knowledge. The weight of the distinction between knowledge and understanding on this view thus falls on the second claim, which merits further consideration.

This thesis is meant to reflect, I think, the internalist dimension to understanding. That is, the idea is that just as one's justification, internalistically conceived, is not undermined by epistemic luck (just the sufficiency of that justification, with true belief, for knowledge), so one's understanding is not undermined either. Closer inspection of this claim reveals that the relationship between understanding and epistemic luck is, however, more complex than Kvanvig and Zagzebski suppose.

The example that Kvanvig offers to illustrate this claim is that of someone who, by reading a book on the Comanche tribe, gains a series of beliefs about the Comanche and, thereby, an "historical understanding of the Comanche dominance of the southern plains of North America from the late seventeenth until the late nineteenth century" (Kvanvig 2003: 197).¹⁸ We are told that the relevant class of beliefs contains no falsehood, and that the agent can answer all the

¹⁸ Understanding of this very general claim might start to look dangerously close to holistic understanding of a subject matter, rather than the non-holistic understanding that we are interested in here. In what follows, I will set this concern to one side and simply read it as non-holistic understanding.

relevant questions correctly in this regard (thereby illustrating that the putative knowledge possessed is not 'opaque'). However, Kvanvig argues that although in such a case one would expect the agent to have knowledge of the relevant beliefs, this is not essential – it could well be, as he points out, that the true beliefs in question have been 'gettierized', perhaps because the information that the agent has is only "accidentally true" (*ibid.*).

I think that a case like this is crucially ambiguous, but we can get a better handle on what is going on here by taking a simpler case and then returning to consider this more complex example in the light of our intuitions as regards the simpler case.

Consider the example of understanding why one's house burned down. Suppose first that we have a standard Gettier-style case in which something intervenes 'betwixt' belief and fact. For example, imagine that, upon finding one's house in flames, one approaches someone who looks as if she is the fire officer in charge and asks her what the reason for the fire is. Suppose one is told by this person that the reason why one's house is burning is faulty wiring, and this coheres with one's wider set of beliefs. But suppose now that the person one asked in this regard is not in fact the fire officer in charge but instead someone who is simply dressed in a fire officer's uniform and who is on her way to a fancy dress party. Still, one did indeed gain a true belief in this regard. So, even though the epistemic luck in question prevents one from having knowledge of the relevant propositions, does one lose one's understanding? Seemingly, it does, for ask yourself the question now of whether you understand why your house burnt down. Surely the answer to this question is a straightforward 'no'. One cannot gain an understanding of why one's house burnt down by consulting someone who, unbeknownst to you, is not the fire officer but instead someone in fancy dress.

12.

So does this mean that Kvanvig is just wrong in thinking that understanding is immune to epistemic luck? Not entirely since, as we have noted above, there is a kind of epistemic luck which is knowledge-undermining but which is not of the sort that appears in Gettier-style cases which intervenes 'betwixt' belief and fact. With this in mind, consider a variant on the case just described in which it is not Gettier-style epistemic luck that is at issue but rather the sort of 'environmental' epistemic luck at issue in the barn facade case. For example, imagine that the apparent fire officer that one asks

about the cause of the fire is indeed the fire officer, but that one could nevertheless have been easily wrong because there were other people in the vicinity dressed as fire officers – all going to the same fancy dress party, say – who one could very easily have asked and who would have given one a false answer (while failing to indicate that they were not real fire officers).

In such a case, as we saw above, one's cognitive success would be because of one's cognitive abilities, and so would constitute a cognitive achievement, and yet the epistemic luck at issue would prevent it from counting as knowledge. The critical question for us, however, is whether it is a case of understanding. I want to argue that it is, and thus that Kvanvig is right on at least this score: 'environmental' epistemic luck, unlike Gettier-style epistemic luck, *is* compatible with possessing understanding. After all, the agent concerned has all the true beliefs required for understanding why his house burned down, and also acquired this understanding in the right fashion. It is thus hard to see why the mere presence of 'environmental' epistemic luck should deprive the agent of understanding.

With this distinction between two types of epistemic luck in mind – one, the Gettier-style epistemic luck, which is inconsistent with understanding and a second, the 'environmental' epistemic luck, which is consistent with understanding – we can return to evaluate Kvanvig's 'Comanche' case. Whether or not the agent retains her understanding in this case will depend on the type of epistemic luck at issue.

So, for example, suppose that the agent forms her beliefs about the Comanches by reading an apparently scholarly book which is in fact nothing of the sort. Let us say, for instance, that the author of this book simply took lots of rumours and unchecked stories about the Comanche and presented them, along with some inventive guesswork, as established fact. But suppose further that despite this lack of attention to scholarship, the author did get matters entirely right. This would thus be a Gettier-style case in which our agent gains lots of true beliefs about the Comanches: she has good reason to think that her beliefs about the Comanche are true, and they are true, but it is just a matter of luck that they are true given that the source of these beliefs is so unreliable. Can one gain an understanding of the Comanche tribe in this way? In particular, can one gain an historical understanding of why the Comanche were so dominant in the southern plains of North America from the late seventeenth until the late nineteenth century in this fashion? I want to suggest that one cannot, any more than one can gain an understanding of why one's house

burnt down by gaining a true belief about what caused the fire from someone pretending to be a fire officer.

Matters are different, however, if we redescribe the case as a form of environmental epistemic luck, rather than as Gettier-style epistemic luck. Suppose, for example, that the book that the agent consults is indeed appropriately scholarly – and thus reliable – when it comes to this subject-matter, and that the agent accordingly gains lots of true beliefs about the Comanche. Nevertheless, the luck enters the picture because of how all the other books on this topic – which are also superficially just as a scholarly – are very unreliable, and one could very easily have found out what one did by consulting one of these books. Does epistemic luck of this sort undermine one's understanding in the way that it would undermine one's knowledge? I don't think that it does, since one did indeed find out the relevant facts in the right kind of way. Just as one can gain an understanding of why one's house burnt down by speaking to the fire officer – even though one could just have easily been misled by someone who isn't the fire officer – so one can gain an understanding of the Comanche by reading a reliable book even though one could have very easily consulted an unreliable book.¹⁹

13.

So while Kvanvig and others are right to think that understanding is compatible with a certain type of knowledge-undermining epistemic luck, they are wrong to think that it is compatible with all types of knowledge-undermining epistemic luck. Their mistake, it seems, is to fail to distinguish between two crucial ways in which epistemic luck can be knowledge-undermining. That understanding is compatible with one type of knowledge-undermining epistemic luck suffices, however, to show that knowledge is distinct from understanding, since it entails that one can have understanding without the associated knowledge.

¹⁹ While noticing that Kvanvig's claim that understanding is compatible with epistemic luck is not quite right, Grimm 2006 fails to recognise that the mistake here is simply to equate environmental epistemic luck with Gettier-style epistemic luck. As a result, he concludes that understanding is just as incompatible with epistemic luck as knowledge is, and thus that knowledge is a species of understanding after all.

One consequence of this is that the standard view of understanding outside of epistemology, such that understanding is a species of knowledge, is false. Indeed, this is not the only respect in which this conception of understanding is mistaken. Recall that on this conception of understanding, to understand why such-and-such is the case is equivalent to knowing why such-and-such is the case, which is in turn equivalent to knowing that such-and-such is the case because of such-and-such. As we have seen, however, the problem of environmental epistemic luck illustrates that I can understand why my house burned down even while failing to know why it burned down (indeed, even while failing to know that it burned down because of faulty wiring).

There is also a second respect in which this conception of understanding is mistaken, since it is possible to know why one's house has burned down (and indeed know that it burned down because of faulty wiring), even though one does not understand why one's house burned down. We can illustrate this point via a Lackey-style example. Suppose that I understand why my house burned down, know why it burned down, and also know that it burned down because of faulty wiring. Imagine further that my young son asks me why his house burned down and I tell him. He has no conception of how faulty wiring might cause a fire, so we could hardly imagine that merely knowing this much suffices to afford him understanding of why his house burned down. Nevertheless, he surely does know that his house burned down because of faulty wiring, and thus also knows why his house burned down. Indeed, we can imagine a teacher asking my son if he knows why his house burned down and him telling the teacher the reason. If asked by a second teacher if my son knew why his house burned down, we could then imagine the first teacher saying that he did. So, it seems, one can not only have understanding without the corresponding knowledge, but also knowledge without the corresponding understanding.²⁰

14.

Just as the Lackey-style case offered earlier demonstrated that sometimes one might have knowledge without a cognitive achievement, the same moral can be drawn here. My son might know why his

²⁰ For more on the relationship between understanding and knowing-why, see Pritchard *forthcomingc*.

house burned down, but this knowledge does not constitute a cognitive achievement on his part because of how he is unable to take appropriate credit for the truth of his belief. Interestingly, however, we have just seen that while knowledge and cognitive achievement come apart on this score, understanding and cognitive achievement do not. My son's knowledge does not constitute a cognitive achievement, but then neither does it constitute genuine understanding on his part either.²¹

Indeed, we have good reason to think that all understanding involves cognitive achievement. Recall that the moral of the barn facade case described earlier was that one could exhibit a cognitive achievement and yet lack knowledge, because of how knowledge, unlike cognitive achievement, is incompatible with environmental epistemic luck. The same applies to understanding. When one couples this observation with the fact that the cases in which an agent has knowledge while not exhibiting a cognitive achievement are cases in which the agent lacks the relevant understanding, then one can see that there is a strong *prima facie* case for thinking that all understanding involves a cognitive achievement.

Indeed, I think this thesis is highly plausible. Its plausibility relates to the fact that understanding seems to be essentially an epistemically internalist notion, in the sense that if one has understanding then it should not be opaque to one that one has this understanding – in particular, one should have good reflectively accessible grounds in support of the relevant beliefs that undergird that understanding. But given that this is a requirement of understanding, it is unsurprising that one can construct a Lackey-style case in which an agent has knowledge but not understanding, since such cases work precisely by using examples of agents who, while having knowledge, lack good reflectively accessible grounds in favour of their beliefs.

That understanding is both factive and resistant to Gettier-style epistemic luck also demonstrates, however, that we should be wary of construing understanding along purely internalist lines. One's reflectively accessible grounds in favour of one's belief might well survive the falsity of what one believes and also be compatible with Gettier-style luck, but as we have seen, the same is not true of understanding. Just as genuine cognitive achievements do not depend exclusively on the cognitive efforts of the agent, but also on the relevant cognitive success and the right connection obtaining between

²¹ I argue in Pritchard *forthcominga* that this point has some important implications for the epistemology of testimony.

cognitive ability and cognitive success, so genuine understanding makes the same 'external' demands.

15.

So where does all this leave us as regards the problem of epistemic value? Recall that we noted above that the robust virtue epistemic approach to this problem did not succeed, in that there was no straightforward way of showing that the distinctive value of cognitive achievements carried over to knowledge. At most, this approach demonstrated that it is cognitive achievements that are distinctively valuable but, since one could exhibit a cognitive achievement while lacking knowledge, and know while failing to exhibit a cognitive achievement, this thesis, unless suitably supplemented with further argument at any rate, did not translate into the claim that knowledge is distinctively valuable. Indeed, as matters stand, what we end up with is a kind of revisionism rather than a form of validationism, in that it is actually cognitive achievements that are distinctively valuable, rather than knowledge.

We have noted here that understanding also comes apart from knowledge, in the sense that one can have understanding while lacking the corresponding knowledge and have knowledge while lacking the corresponding understanding. Nevertheless, if cognitive achievements are the kind of thing which are distinctively valuable, then – given that we have seen that there are strong grounds for supposing that understanding is a kind of cognitive achievement – we have a straightforward explanation of why understanding is distinctively valuable.

At the very least, then, we have a form of revisionism available to us which could explain why understanding, rather than knowledge, is distinctively valuable. Interestingly for our purposes, however, a form of revisionism which appeals to the special value of understanding is, I think, more appealing than a form of revisionism which appeals to an epistemic standing which is clearly a 'lesser' epistemic standing when assessed relative to knowledge, such as justified true belief. For while one can have understanding while lacking knowledge, it should be clear that understanding requires an intellectual sophistication that is not necessarily demanded by knowledge. One can imagine, for example, an agent knowing a great deal while having very little understanding of anything, but it is hard to imagine the converse. If understanding is a cognitive achievement, something that can only fall short of knowledge when environmental

epistemic luck is present, then we have a straightforward explanation for this intuition. Knowledge may or may not be distinctively valuable, but understanding certainly is, and given the features of understanding that we have noted, this claim, while in itself revisionist, is certainly highly plausible.²²

Edinburgh University

References

- Achinstein, P. (1983) *The Nature of Explanation*, Oxford University Press, Oxford.
- Brogaard, B. (2007) "I Know. Therefore, I Understand", *typescript*.
- Chisholm, R. (1977) *Theory of Knowledge* (2nd Ed.), Prentice-Hall, Englewood Cliffs, New Jersey.
- Elgin, C. (1996) *Considered Judgement*, Princeton University Press, Princeton, NJ.
- (2004) "True Enough", *Philosophical Issues* **14**, 113–31.
- (Forthcoming). "Is Understanding Factive?", *Epistemic Value*, (eds.) A. Haddock, A. Millar & D. H. Pritchard, 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, UK.
- (2002) "Knowledge as Credit for True Belief", *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) M. DePaul & L. Zagzebski, Oxford University Press, Oxford.
- (2007) "The Nature of Ability and the Purpose of Knowledge", *typescript*.

²² An earlier version of this paper was presented at the *Royal Institute of Philosophy* as part of the *Epistemology* lecture series in 2006 and at a departmental talk at the University of Edinburgh in 2007. I am grateful to the audiences on these occasions, especially Sharar Ali, Matthew Chrisman, Andy Clark, Jennifer Hornsby, Jesper Kallestrup, Anthony O'Hear, Mike Ridge and Barry Smith. Special thanks go to Ernie Sosa, who commented on an earlier draft. This paper has benefited from an AHRC Research Leave award.

- (Forthcominga) “The Value Problem”, in *The Value of Knowledge*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, Oxford University Press, Oxford.
- (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**, 515–35.
- Kitcher, P. (2002) “Scientific Knowledge”, in *Oxford Handbook of Epistemology*, (ed.) P. Moser, Oxford University Press, Oxford.
- Kvanvig, J. (2003) *The Value of Knowledge and the Pursuit of Understanding*, Cambridge University Press, Cambridge, UK.
- (Forthcoming) “Responses”, in *The Value of Knowledge*, (eds.) A. Haddock, A. Millar, & D. H. Pritchard, Oxford University Press, Oxford.
- Lackey, J. (2007) “Why We Don’t Deserve Credit for Everything We Know”, *Synthese* **156**.
- Lipton, P. (2004) *Inference to the Best Explanation*, Routledge, London.
- Pritchard, D. H. (2002). “Resurrecting the Moorean Response to the Sceptic”, *International Journal of Philosophical Studies* **10**, 283–307.
- (2005) *Epistemic Luck*, Oxford University Press, Oxford.
- (2006) “Knowledge, Luck and Lotteries”, in *New Waves in Epistemology*, (eds.) V. F. Hendricks & D. H. Pritchard, Palgrave Macmillan, London.
- (2007a) “Anti-Luck Epistemology”, *Synthese* **156**.
- (2007b) “Recent Work on Epistemic Value”, *American Philosophical Quarterly*, **44**, 85–110.
- (2007c) “Sosa On Epistemic Value”, 2nd On-Line Philosophy Conference, http://experimentalphilosophy.typepad.com/2nd_annual_online_philoso/2007/05/ernest_sosa.html
- (2007d) “The Value of Knowledge”, *typescript*.
- (Forthcominga). “A Defence of Quasi-Reductionism in the Epistemology of Testimony”, *Philosophica*.
- (Forthcomingb) “Greco on Knowledge: Virtues, Contexts, Achievements”, *The Philosophical Quarterly*.
- (Forthcomingc) “Knowing the Answer, Understanding and Epistemic Value”, *Grazer Philosophische Studien*.
- 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.

Knowledge, Understanding and Epistemic Value

- Riggs, W. (Forthcoming) "Getting the *Meno* Requirement Right", in *Epistemic Value*, (eds.) A. Haddock, A. Millar & D. H. Pritchard, Oxford University Press, Oxford.
- Salmon, W. (1989) "Four Decades of Scientific Explanation", *Minnesota Studies in the Philosophy of Science* 13.
- Sartwell, C. (1992) "Why Knowledge is Merely True Belief", *Journal of Philosophy*, **89**, 167–80.
- Sosa, E. (1988) "Beyond Skepticism, to the Best of our Knowledge", *Mind* **97**, 153–89.
- (1991) *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge University Press, Cambridge, UK.
- (1999) "How to Defeat Opposition to Moore", *Philosophical Perspectives* **13**, 141–54.
- (2007) *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford University Press, Oxford.
- Stich, S. (1990) *The Fragmentation of Reason*, MIT Press, Cambridge, MA.
- Unger, P. (1968) "An Analysis of Factual Knowledge", *Journal of Philosophy* **65**, 157–70.
- Williamson, T. (2000) *Knowledge and its Limits*, Oxford University Press, Oxford.
- Woodward, J. (2003) *Making Things Happen: A Theory of Causal Explanation*, 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?", in *The Blackwell Guide to Epistemology*, (eds.) J. Greco & E. Sosa, 92–116, Blackwell, Oxford.
- (2001) "Recovering Understanding", in *Knowledge, Truth and Duty: Essays on Epistemic Justification, Responsibility and Virtue*, (ed.) M. Steup, Oxford University Press, Oxford.
- (2003) "The Search for the Source of the Epistemic Good", *Metaphilosophy* **34**, 12–28; and reprinted in *Moral and Epistemic Virtues*, (eds.) M. S. Brady & D. H. Pritchard, 13–28, Blackwell, Oxford (2003).