



Kelp, C. (2014) Epistemology extended. *Philosophical Issues*, 24, pp. 230-252.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/140960/>

Deposited on: 16 May 2017

Enlighten – Research publications by members of the University of Glasgow
<http://eprints.gla.ac.uk>

Epistemology Extended

CHRISTOPH KELP*

Abstract

A common presupposition in epistemology is that the processes contributing to the generation of knowledge do not extend beyond the knower's skin. This paper challenges this presupposition. I adduce a novel kind case that causes trouble for a number of even the most promising accounts of knowledge in current literature (virtue epistemological and modal accounts), at least so long as the presupposition is in place. I then look at a couple of recent accounts of knowledge that drop the presupposition and expressly allow the relevant processes to extend beyond the knower's skin. While these accounts can handle the problem case, they encounter difficulties elsewhere: extension occurs too easily and so the accounts predict knowledge where they ought not. Finally, I offer a novel way of extending epistemology and argue that it can steer clear of the problems on both sides.

Introduction

What is knowledge? It was once widely believed that the answer is that knowledge is justified true belief. Thanks to Gettier [1963], we know that this can't be right: one can have a justified true belief and yet fail to know. We also know that there is no quick way to fix this shortcoming of the JTB account. The post-Gettier literature teaches us as much.

That said, there are a number of promising accounts of knowledge in recent literature. One of them has been advanced by virtue

*Centre for Logic and Analytic Philosophy, KU Leuven, Kardinaal Mercierplein 2, BE-3000 Leuven, Email: christoph.kelp@hiw.kuleuven.be

epistemologists.¹ The idea here is that knowledge is a kind of success due to ability or competence. According to another promising proposal, knowledge features a modal condition. The two most prominent candidates in this camp are that knowledge requires belief that is sensitive to falsehood² or safe from error³.

This paper aims to develop a novel type of problem for both virtue epistemological and modal accounts of knowledge, at least as they are traditionally understood. Crucially, I do not intend to provide another Gettier-style counterexample, i.e. a type of case that shows that the conditions suggested by these accounts are not sufficient for knowledge after all. Rather, I will try to describe a type of case that suggests that the conditions imposed are not necessary, i.e. a case of knowledge in which the suggested conditions are not met (§§2,3).

As I see it, what we will be dealing with is a case of extended cognition (EC), in a sense that, I hope, will become clear in due course. That is why, with the argument against traditional versions of virtue epistemological and modal accounts of knowledge in place, I will look at a number of ‘extended epistemologies’, which combine core ideas of relevant virtue epistemological and/or modal accounts of knowledge with ideas from the literature on extended cognition. I will argue that two existing extended epistemologies remain ultimately unsatisfactory and develop a novel way of extending epistemology that is more promising (§4).

The idea that cases of extended cognition cause trouble for at least a certain type of virtue epistemological (VE) account of knowledge is not new. It has been defended by Duncan Pritchard and Krist Vaesen in a number of recent pieces.⁴ I do not think that their arguments are ultimately successful. There is a kind of virtue epistemological

¹ Virtue epistemology has been defended, among others, by John Greco [2003, 2010, 2012], Ernest Sosa [2007, 2010, 2011], Wayne Riggs [2002, 2009a, 2009b] and Linda Zagzebski [1996]. I have also defended a couple of versions of the view in [Kelp 2011, 2013d].

² Champions of sensitivity accounts include Robert Nozick [1981], Fred Dretske [1971, 1981] and more recently, Tim Black and Peter Murphy [2007, 2008] and Kelly Becker [2007, 2009].

³ For defences of safety see [Sainsbury 1997, Sosa 1999, Williamson 2000, Pritchard 2005]. For an account that combines a safety condition and a virtue epistemological condition see [Pritchard et al. 2010, Pritchard 2012, Kelp 2013c].

⁴ [Pritchard 2010, Vaesen 2011, 2013]. It may be worth noting that Vaesen’s counterexamples involve at best a very weak version of extended cognition and that it remains an open question whether his examples are not best seen as involving cases of embedded cognition. (Thanks to Orestis Palermos for pointing this out to me.)

account of knowledge that escapes them unscathed. While I have already argued this point elsewhere [Kelp 2013a,b] in some detail, it will be instructive to rehearse some of the relevant arguments here, if only to bring the contribution of this paper into proper focus. I will tackle this task in §1.

1 Pritchard's and Vaesen's arguments

According to the type of VE Pritchard and Vaesen attack, knowledge is a kind of success due to ability. More specifically, in the case of knowledge the relevant kind of success is true belief and the relevant kind of ability is a cognitive ability, that is, roughly, a disposition to form true beliefs. While champions of the kind of VE Pritchard and Vaesen attack typically hold that their preferred virtue condition on knowledge is not only necessary but also sufficient for knowledge, the following weaker claims will do for the purposes of this paper:

VE₁. *S* knows that *p* only if *S*'s believing *p* truly is due to the exercise of cognitive ability.

VE₂. *S*'s belief that *p* is degettiered if and only if *S*'s believing *p* truly is due to the exercise of cognitive ability.⁵

A core question that champions of VE will have to answer is what it takes for a success to be 'due to' the exercise of ability. There are two main contenders in the literature. The first, associated with John Greco [2003, 2010, 2012], unpacks 'due to' in terms of creditability. Roughly, the idea is that a success is due to the exercise of ability just in case it is creditable to the exercise of ability. In contrast, the second, which is most prominently championed by Ernest Sosa [2007, 2010, 2011], unpacks 'due to' in terms of competence manifestation. Here

⁵ Let's distinguish with Hetherington [e.g. 1999] and Pritchard [e.g. 2008] between two types of Gettier case. The first one comprises cases in which agents acquire their beliefs in ways that are epistemically defective in some way (e.g. via inferences from false premises, deceptive experiences, taking readings from faulty equipment, etc.). In contrast, the second type of Gettier comprises cases in which agents do not acquire their beliefs in epistemically defective ways but might easily have done so. For the purposes of this paper, we can restrict the range of Gettier cases that VE₂ claims to account for to cases of the first type. Since, as even champions of VE acknowledge, Gettier cases of the second type pose a serious challenge to their view, it is of some significance that this paper's argument can proceed under this restriction. It means that we can be certain that the problem does not reduce to the already well documented difficulty that arises from cases of the second type.

the rough idea is that a success is due to the exercise of ability just in case it manifests ability.

1.1 Creditability VE

Vaesen and Pritchard's arguments focus on versions of VE that unpack 'due to' in terms of creditability. The question its champions will need to answer is to what degree a given success must be creditable to the exercise of ability. Pritchard and Vaesen argue that it must be *primarily* so creditable. Otherwise, VE will be incapable of explaining the absence of knowledge in Gettier cases [Pritchard 2010: 137, Vaesen 2011: 518]. The problem is that this leaves VE vulnerable to another problem: there are cases in which the agents have knowledge, but their cognitive successes, if creditable to the agents at all, are not primarily so creditable. The literature features a number of different candidates for such cases, the most prominent ones being certain cases of testimonial knowledge [e.g. Lackey 2007, 2009, Pritchard 2008, Pritchard et al. 2010]. Crucially for present purposes, Vaesen and Pritchard argue that cases of extended cognition also fit the bill. Vaesen even goes so far as to suggest that certain cases of extended cognition may well be the best such cases [Vaesen 2011: 522].

Pritchard and Vaesen offer different cases of extended cognition in order to argue against VE. Since I have discussed Vaesen's case at length elsewhere [Kelp 2013b], I will here focus on Pritchard's case:

Notebook. In order to counter his Alzheimer's-induced loss of memory, Otto's wife sets Otto up with a notebook in which Otto is to store and from which he is to retrieve information in much the same way in which he had hitherto stored information in and retrieved it again from memory. Otto's wife teaches Otto how to use the notebook and plays a crucial part in ensuring that the system functions effectively. In virtue of it being an entry in his notebook, Otto believes (or comes to believe now that he retrieves the information) that he has an appointment with the doctor today [Pritchard 2010: 149].

Pritchard points out that, intuitively, Otto knows that he has an appointment with the doctor. At the same time, his true belief does not appear to be primarily creditable to the cognitive abilities that Otto exercises here. Otto's wife gets too large a share of the credit for this to be plausible.

If this is correct, VE faces a dilemma. Either the degree of creditability required for knowledge is construed robustly in terms of pri-

mary creditability. In that case, VE may be able to correctly predict lack of knowledge in Gettier cases, but it will also, and now mistakenly, predict lack of knowledge in cases like *Notebook*. Or the degree of creditability required for knowledge is construed weakly as not requiring primary creditability. In that case, VE can correctly predict knowledge in the latter cases, but will succumb to Gettier cases.

I remain unconvinced by a number of claims assumed by Pritchard and/or Vaesen, including (i) that VE must unpack 'due to' in terms of primary creditability in order to avoid Gettier cases and (ii) that cases of extended cognition may well be the best cases to make the point (Vaesen). That said, I also believe that the argument against creditability VE is successful independently of whether (i) and (ii) hold.⁶

1.2 Competence Manifestation VE

Even if we grant that Pritchard and Vaesen identify a genuine problem for creditability VE, the question remains whether their argument is equally successful against Sosa's competence manifestation version of the view. As I will argue in the following paragraphs, the answer to this question is 'no'. Competence manifestation VE can escape their arguments unscathed. I will again focus on Pritchard and refer the reader to my [Kelp 2013b] for a detailed discussion of Vaesen's case.

In order to prepare the ground for my response to Pritchard, it will first be necessary to briefly sketch the most important parts of Sosa's [2010, 2011] account of competences and competence manifestation. According to Sosa, competences are dispositions to perform well. Sosa distinguishes three components of dispositions (and hence of competences), to wit, constitutional (= CO), inner (= IN) and situational (= SI). In the case of archery competence, the CO are certain properties of the agent, including more basic motor and cognitive competences, the IN include being awake, attentive, etc. and the SI include there being enough light, normal winds, etc. [Sosa 2010: 465].

Corresponding to these components of competences are three levels of competence. The first, 'constitutive' competence includes only the CO, the second, 'inner' competence includes the IN besides the CO, and the third, 'complete' competence includes the CO, IN and SI. When Sosa says that competences are dispositions to perform

⁶ For a more detailed argument that even works for Pritchard's weakened creditability version of VE, see [Kelp 2013c].

well what he means is that complete competences are dispositions to perform well in the sense that when an agent has the complete competence—i.e. CO, IN and SI are in place—were he to produce a relevant performance it would highly likely be successful. For an archer to have the competence to hit the target is for him to have a disposition to hit the target in the sense that if the agent were to produce a shot in suitable CO, IN and SI, his shot would very likely hit the target [Sosa 2010: 466].

Crucially, according to Sosa, a success manifests competence only if it manifests a *complete* competence [Sosa 2010: 470]. That is to say, the competence's CO, IN and SI must all be in place. When a competence is exercised in unsuitable SI, the performance it produces won't qualify as a success that manifests competence, not even if it turns out to be successful. For instance, when an archer's shot is blown off target by a gust of wind, the competence's SI are not satisfied. As a result, his shot won't qualify as a success that manifests competence, not even if it does after all hit the target, say because a second gust of wind brings it back on target.

Recall that, according to Sosa's version of VE, the cognitive success at issue in knowledge is *due to* the exercise of cognitive competence just in case it *manifests* competence. With the above account of competences and competence manifestation in play, we now know that this means that one's cognitive success must manifest a *complete* competence. This serves to explain why beliefs of agents in standard Gettier cases do not qualify as knowledge. In such cases, the relevant competence's SI are not satisfied (as e.g. the agent is looking at a surface that is illuminated by coloured light, at a non-sheep that looks just like a sheep, at a stopped clock, etc.). As a result, in standard Gettier cases, the agent's cognitive success does not manifest competence. Sosa's account predicts, correctly, that agents in such cases lack knowledge.

Of course, in order to escape Pritchard and Vaesen's argument, it remains to be shown that the present account of competence manifestation can allow that the true beliefs in cases like *Notebook* do manifest competence. To see this, I will first have to add a few more details to Sosa's account of competences.

Recall that, according to Sosa, competences are dispositions to perform well. While this is roughly right, some cases suggest that a slight refinement is required. Suppose an agent has more than one way of producing performances of a certain kind. To keep things

simple, let's say he has two such ways. For instance, our archer may shoot with his left hand or with his right hand. Suppose he is disposed to perform well when producing performances of that kind in one way but not the other. Call the former 'the good way' and the latter 'the bad way'. Our archer may be disposed to hit the target when shooting with his right hand, but not when shooting with his left hand. Suppose, finally, that the agent is disposed to produce performances of said kind in the bad way. His dominant way of producing performances is the bad way. Our archer may be disposed to shoot with his left hand. In this situation, there is a clear sense in which the agent is not disposed to perform well. Even so, the agent possesses a competence to achieve the result, *viz.* when producing performances in the good way. It's just that he is disposed not to exercise this competence. These considerations suggest that competences are relative to ways of producing performances. Competences are dispositions to perform well via certain ways of performance production.

Notice next that ways of producing performances may involve equipment. Our archer's way of producing shots involves a bow and arrow. Note also that the SI of equipment-involving competences may include conditions concerning the quality of the equipment. Our archer may have the competence to hit the target only in SI such that his set of bow and arrows are in sufficiently good state. He does not have the competence to hit the target when the quality of his bow is degraded to the point of dysfunctionality.

Now, it may be that the (continued) satisfaction of a competence's SI is ensured by the efforts of another agent, a supporter. Suppose, for instance, our archer owns two bows. One bow is now degraded to the point of dysfunctionality. The other is more durable, at least given that it is suitably maintained. Suppose the durable set has been designed and is now maintained by another agent, a supporter. The supporter also sees to it that our archer uses the durable set when shooting. Here the supporter ensures that (some of) the SI of our archer's competence are in place. What's more, we can easily describe the case in such a way that had it not been for the efforts of the supporter, our archer would now be in unsuitable SI and that this unfortunate situation might very easily have obtained. In that case, part of the credit for any success our archer goes on to attain will go to his supporter. Perhaps the supporter's share in the credit is even large enough that these successes are no longer primarily creditable to our archer. However, and this is the crucial point, it

does not change the fact that, when SI (as well as CO and IN) are suitable, our archer rises to the level of complete competence and so, by Sosa's lights, may produce successes that manifest competence. By Sosa's lights, what is going on in such cases is that our archer rises to the level of complete competence, although he might very easily have failed to do so. Again, while this may mean that the degree of credit he deserves is diminished, it does not mean that, now that he does rise to the level of complete competence, his successes do not manifest competence.

The situation is essentially the same in *Notebook*. Here Otto has two types of memory, as it were, his biological memory and his external memory, the notebook. His biological memory is degraded to the point of dysfunctionality. The notebook is more durable, at least when suitably maintained. It was designed and is now maintained by Otto's wife who also sees to it that Otto uses it when forming memory based beliefs. Now consider Otto's true belief, formed in a way involving his notebook in otherwise suitable SI (as well as CO and IN), that he has an appointment with the doctor today. That Otto believes truly may be creditable to a significant degree to his wife. It may be creditable to her in such a way that it can no longer be said to be primarily creditable to Otto. However, this is not to say that Otto's believing the truth does not manifest his competence to form true beliefs, now that he is forming beliefs in suitable SI. On the contrary, by Sosa's lights, it does.

Sosa's competence manifestation version of VE can accommodate both the intuition of ignorance in standard Gettier cases and the intuition of knowledge in cases like *Notebook*. Competence manifestation VE escapes Pritchard's and Vaesen's attack unscathed.

2 The Timekeeper Case

While Pritchard and Vaesen's arguments remain unsuccessful at least against competence manifestation VE, this section develops a more serious problem for the view. To get there let's first look at a familiar case:

The Original. The timeseeker looks at a public clock, sees that it reads 2.30 and on that basis comes to believe that it is 2.30. The clock has an outstanding track-record of functioning properly and the timeseeker has no reason to think that it is currently not accurate. Her belief is true. It is in fact 2.30. Unbeknownst to the timeseeker, however, the clock has stopped exactly twelve hours ago.

Intuitively, the timeseeker does not know that it's 2.30. Competence manifestation VE can readily accommodate this intuition. Among the SI of the timeseeker's clock-reading competence is that the clock is functioning properly. In *The Original*, this condition is not satisfied. For that reason, competence manifestation VE predicts that the timeseeker's true belief does not manifest his clock-reading competence and therefore falls short of knowledge.

Now consider the following variation of the case:

The Timekeeper. The timeseeker looks at a public clock, sees that it reads 2.30 and on that basis comes to believe that it is 2.30. The clock has an outstanding track-record of functioning properly and the timeseeker has no reason to think that it is currently not accurate. Her belief is true. It is in fact 2.30. Unbeknownst to the timeseeker, however, the clock has stopped exactly twelve hours ago. As it happens, this episode is observed by the timekeeper, who has been called in to fix the stopped clock. Using his two radio clocks, the timekeeper confirms that the reading of the stopped clock is accurate. Had the stopped clock reading been inaccurate, the timekeeper would have alerted the timeseeker to this fact.

Here is my intuition about the case: The timeseeker *knows* that it's 2.30.⁷ Competence manifestation VE will find it very difficult to accommodate this intuition, at least whilst continuing to secure the right verdict in *The Original*. Here is why. In both cases the timeseeker acquires his belief in the same way, via an exercise of his clock-reading competence. As we have already seen, it is plausible that the SI of the timeseeker's clock-reading competence include that the clock be functioning properly. Since in the present case, this condition is not satisfied, competence manifestation VE will predict that, here too, the timeseeker's true belief does not manifest competence and therefore falls short of knowledge.⁸

⁷ If you do not share this intuition, consider a variation of the case in which the timekeeper's involvement is modally robust rather than a mere accident. Notice that if you now get the intuition of knowledge, the case will still serve to generate a problem for both non-extended versions of VE and for (VE combined with) the extended epistemologies discussed in §§4.1, 4.2. (Thanks to Jesper Kallestrup for drawing my attention to this.) If you still do not have the intuition that the timeseeker knows, I will provide some independent reason to think that he does in the conclusion, at least for those who accept the hypothesis of extended cognition.

⁸ *The Timekeeper* is no less problematic for creditability VE. After all, it would seem that the degree to which the timeseeker's cognitive success is creditable to her is the same in *The Original* and *The Timekeeper*. If so, creditability VE will be

3 Modal Accounts of Knowledge

It turns out that VE is in trouble. This might be an opening for champions of modal accounts of knowledge. If some modal account can avoid the problems that VE encounters, we might have reason to favour it over VE. For that reason, in what follows, I will look at the two most prominent modal accounts of knowledge and investigate whether they can successfully handle the two versions of the stopped clock case.

3.1 Sensitivity

According to the first of these accounts, knowledge requires belief that is sensitive to falsehood. Roughly, this means that one knows that p only if were p false, one would not believe that p . While this is indeed roughly right, it is now widely acknowledged that a further refinement is needed. The sensitivity condition must be indexed to ways (or methods) of belief formation. Here is how Nozick, one of the most prominent advocates of this kind of view, formulates his version of the sensitivity condition (SEN):

SEN₁. S knows that p via method M only if: “if p weren’t true and S were to use M to arrive at a belief whether (or not) p , then S wouldn’t believe, via M , that p .” [Nozick 1981: 179]

In addition, I will, again in line with Nozick, assume a standard possible worlds semantics for counterfactual conditionals of the kind at issue in the sensitivity condition. According to this semantics, a sentence of the form “if ϕ were the case, then ψ would be the case” is true just in case ψ is true at the closest worlds at which ϕ is true. This gives us the following alternative statement of the sensitivity account:

SEN₁. S knows that p via method M only if: at all the closest worlds at which p is not true and S uses M to arrive at a belief whether (or not) p , then S does not believe, via M , that p .

It is worth noting that Nozick countenances a further condition on knowledge besides SEN that is meant to do some anti-Gettier duty. Since all the Gettier cases in this paper are either handled by SEN₁ or else satisfy Nozick’s further condition, for present purposes it will be safe to assume the following additional thesis:

hard pressed to accommodate both the intuition of ignorance in the former and the intuition of knowledge in the latter.

SEN₂. *S*'s belief that *p* via method *M* is degettiered if and only if: at all the closest worlds at which *p* is not true and *S* uses *M* to arrive at a belief whether (or not) *p*, then *S* does not believe, via *M*, that *p*.

Just as VE, SEN handles *The Original* without further difficulties. Among the closest worlds at which it is not 2.30 and the timeseeker forms a belief that it is 2.30 via his clock-reading competence, are worlds at which the clock that she looks at continues to be stopped at 2.30. At those worlds, the timeseeker will form a false belief that it is 2.30 via her clock-reading competence. So SEN is not satisfied. SEN predicts, correctly, that the timeseeker does not know.

But what about *The Timekeeper*? Does SEN allow us to make any progress on this front? One might initially think that it does. After all, at the closest worlds at which it is not 2.30, it is 2.29 or 2.31. At those worlds, one might think, the timekeeper would step in and inform the timeseeker that it is in fact 2.29 rather than 2.30 in which case the timeseeker would not believe that it is 2.30 via his clock-reading competence. SEN appears to be satisfied.

On second thought, however, it is not so clear that this is right. To see this, let's ask what else happens at the closest worlds at which it is not 2.30. Does the timekeeper continue to arrive at 2.30 in at least some of them? If so, then there will be some subset of the closest worlds at which the timeseeker will already have formed his false belief that it's 2.30 before the timeseeker has a chance to correct him. This problem would be avoided, if worlds at which the timekeeper continues to watch over the timeseeker are closer than worlds at which he continues to arrive at 2.30. Unfortunately, it is hard to see on what grounds it could be argued that the former worlds are closer than the latter. What's even worse, the case can easily be redescribed in such a way that it is plausible that worlds at which the timekeeper watches over the timeseeker are further away than worlds at which he does not watch over him at all. All that needs to be done is add that, by some accident, the timeseeker was delayed and for that reason arrived at the clock later than she otherwise would have, while, again due to some accident, the timekeeper was lucky enough to arrive at the clock earlier than he otherwise would have. In that case, it would seem that the closest worlds at which it is not 2.30 include worlds at which the timeseeker and the timekeeper do not overlap. At those worlds, the timeseeker will of course form a false belief that it is 2.30 via his clock-reading competence. SEN predicts, incorrectly, that the timeseeker's actual belief falls short of knowledge.

3.2 Safety

For many the failure of SEN will not come as a big surprise. After all, the view faces a number of problems that many would consider even more serious than the one that may arise from making the wrong prediction in a case like *The Timekeeper*. Indeed problems besetting SEN have motivated sympathisers with modal accounts to prefer a safety-based account of knowledge (SAF), which is often thought to fare better than SEN. According to a rough version of SAF, knowledge requires that one might not easily have been in error. Again, the rough version is refined by an index to ways of belief formation. The resulting version of SAF can be stated as follows:

SAF₁. *S* knows that *p* via method *M* only if: not easily might *S* arrive at a false belief on whether *p* via *M*.

SAF₂. *S*'s belief that *p* via method *M* is degettiered if and only if: not easily might *S* arrive at a false belief on whether *p* via *M*.

Given a standard possible worlds semantics for 'might easily' according to which a sentence of the form "It might easily be the case that ϕ " is true if and only if at some close possible world ϕ is true, we get the following alternative version of the safety condition:

SAF₁. *S* knows that *p* via method *M* only if: there are no close possible worlds such that *S* arrives at a false belief on whether *p* via *M*.

SAF₂. *S*'s belief that *p* via method *M* is degettiered if and only if: there are no close possible worlds such that *S* arrives at a false belief on whether *p* via *M*.

Does SAF fare any better than SEN? No. There are many close possible worlds at which the timeseeker arrives a minute early whilst the timekeeper continues to arrive at the same time. At those possible worlds, the timeseeker acquires a false belief that it is 8.30 via his clock-reading competence. Her belief is thus not safe. SAF also makes the incorrect prediction that the timeseeker does not know.⁹

⁹ Couldn't champions of a safety condition venture to avoid modify the requirement imposed on knowledge? Pritchard has argued that there is independent reason to opt for the following weaker version of safety:

SAF₁^{*}. *S* knows that *p* via method *M* only if: *S* avoids false belief on whether *p* via *M* at all very close possible worlds and at most close possible worlds [Pritchard 2007: 292].

The Timekeeper thus poses a problem not only for VE but also for modal accounts of knowledge. This is bad news as it means that two of the most promising types of account in recent epistemology, perhaps the two most promising types, are in trouble. The next section investigates how, if at all, we might be able to deal with this problem.

4 Extended Epistemologies

The timekeeper plays a crucial part in *The Timekeeper*. A core reason why the timeseeker knows the time in that case is that the timekeeper (i) independently verifies the time and (ii) would alert the timeseeker had the clock reading been inaccurate. The trouble for standard virtue epistemological and modal accounts of knowledge arises from the fact that they do not make provisions for positive epistemic status's depending on other agents in this way. How can this defect be remedied?

4.1 Green's Joint Abilities Virtue Epistemology

Adam Green [2012] has recently offered an extended version of virtue epistemology that, one might think, holds out the hope of doing the trick for us. In an effort to marry creditability VE with the idea that the processes involved in the generation of testimonial belief include the processes involved in the production of testimony [Goldberg 2010], he proposes the following 'joint abilities virtue epistemology' (JAVE):

JAVE₁. "[S] knows that p only if the abilities that contribute to the formation and sustenance of [S]'s belief that p deserve primary credit (or something close to it) for [S] knowing p *whether those abilities are contributed solely by [S] or also by other agents.*" [Green 2012: 125]

SAF₂^{*}. S 's belief that p via method M is degettiered if and only if: S avoids false belief on whether p via M at all very close possible worlds and at most close possible worlds.

There is little reason to think that this weakening of safety will allow us to make progress on *The Timekeeper*. After all, it is plausible that worlds at which the timeseeker arrives a minute early while the timekeeper continues to arrive at the same time are among the very close possible worlds. If so, SAF^{*} is not satisfied either. What's more, if the case is set up in the way described in the above discussion of SEN, where it is accidental that the timeseeker arrives as late as at 8.30 and the timekeeper as early as 8.30, at most close worlds, the timeseeker arrives early and doesn't overlap with the timekeeper. At those worlds, she will form a false belief that it is 8.30 via her clock-reading competence.

JAVE₂. *S*'s belief that *p* is degettiered if and only if the abilities that contribute to the formation and sustenance of *S*'s belief that *p* deserve primary credit (or something close to it) for *S* knowing *p* whether those abilities are contributed solely by *S* or also by other agents.¹⁰

Initially it may not be obvious just how JAVE could allow us to make progress on *The Timekeeper*. After all, it is not clear that the timekeeper's abilities contribute to the formation and sustenance of the timeseeker's belief that it is 8.30. But, of course, if they don't, JAVE is bound to predict ignorance as well.

That said, there is a way of looking at the case in which the timekeeper's ability does make a contribution to the formation and sustenance of the timeseeker's belief. To see how, notice first that we are equipped with monitoring abilities, which detect unfavourable environmental conditions and lead us to refrain from forming beliefs or to retract beliefs already formed. For instance, upon realising that the only light source of the room is a blacklight, I refrain from forming beliefs about the colours of the objects around me (by colour vision only). But now notice that, in *The Timekeeper*, the timekeeper's abilities play a similar role for the timeseeker's belief about the time. Were environmental conditions unfavourable, i.e. were the clock not to display the right time, the timekeeper would alert the timeseeker of this fact, which, in turn, would lead the latter to refraining from forming a belief. I want to suggest, then, that we may think of the timekeeper as contributing a kind of *monitoring ability* to the formation/sustenance of the timeseeker's belief. It's just that, in the case of the timeseeker and the timekeeper, these abilities are external to the relevant cognitive agent (the timeseeker) rather than internal.

If we view the timekeeper as contributing a monitoring ability to the timeseeker's belief, it would seem that JAVE is ideally suited to pass the correct verdict in *The Timekeeper*. It may be hard to deny that the contribution of the timeseeker's clock-reading ability by itself does not deserve primary credit for her cognitive success. At

¹⁰ Note that Green should have no qualms with JAVE₂. After all, Green explicitly claims that he wants to retain "the spirit of Greco" [Green 2012: 125]. It is not hard to see that JAVE won't do that unless Green is prepared to accept at least JAVE₂ besides JAVE₁. After all, according to Greco, a virtue condition is not only necessary but also *sufficient* for knowledge. Furthermore, there is reason to believe that Green does in fact accept a thesis entailing JAVE₂ when he claims that virtue epistemologies, like Greco's and JAVE, do not merely aim to specify "a necessary condition on knowledge but ... *the* necessary condition on knowledge ... that defines the epistemic enterprise." [Green 2012: 125]

the same time, it is fairly plausible that the joint contribution of the timeseeker's clock-reading ability and the timekeeper's monitoring ability do deserve such primary credit. If so, JAVE can successfully accommodate the intuition of knowledge in *The Timekeeper*.

It is also worth noting that this does not compromise JAVE's ability to explain lack of knowledge in a wide range of standard Gettier cases. Thus consider *The Original* once more. The only abilities that contribute to the formation of the timeseeker's belief are his own abilities. At the same time, the contribution to cognitive success made by the timeseeker's abilities alone is no more deserving of primary credit here than in *The Timekeeper*. In consequence, JAVE can successfully predict lack of knowledge in *The Original*.

While JAVE may thus look promising, unfortunately, there is reason to believe that it remains unsuccessful after all. There is reason to believe that it succumbs to Gettier cases of a different kind. Here is one such case:

The Gettieriser. The timeseeker looks at a public clock, sees that it reads 2.30 and on that basis comes to believe that it is 2.30. The clock has an outstanding track-record of functioning properly and the timeseeker has no reason to think that it is currently not accurate. Her belief is true. It is in fact 2.30. Unbeknownst to the timeseeker, however, the clock is stopped but nonetheless displays the time accurately because the gettieriser, who tries to maximise the number of gettiered belief in the world, happened to pass by and seized the opportunity by setting the clock exactly to the right time.

In this case, the timeseeker's belief is gettiered (and hence does not qualify as knowledge). At the same time, it is hard to see how JAVE could accommodate this datum. To see this, notice first that, in this case, the abilities of the gettieriser make a very clear contribution to the formation and sustenance of the timeseeker's belief that it is 2.30: the gettieriser contributes the reading of the clock. Now, while it is plausible that the contribution of the timeseeker's own ability is not sufficient to be deserving of primary credit for his cognitive success, the joint contribution of the gettieriser's and the timeseeker's ability would appear to be. If it is, JAVE predicts, incorrectly, that the timeseeker's belief is not gettiered.

Notice that, at this stage, we don't even need to claim that JAVE falls prey to gettierisation. Rather, all that we need to generate a problem for JAVE is the more cautious conditional claim that if the joint contribution of the timeseeker and the timekeeper's abilities in

The Timekeeper is deserving of primary credit, then so is the joint contribution of the timeseeker and the gettieriser's abilities in *The Gettieriser*. The conditional claim suffices to generate a familiar type dilemma for JAVE: If the timeseeker and timekeeper's joint contribution is deserving for primary credit in *The Timekeeper*, JAVE can explain the intuition of knowledge here. Unfortunately, however, by the conditional claim, champions of JAVE will also have to acknowledge that the joint contribution of the timeseeker and the gettieriser is deserving of primary credit and JAVE will predict, incorrectly, that the timeseeker's belief in *The Gettieriser* is degettiered. On the other hand, if champions of JAVE were to claim that the contribution of the timeseeker and the gettieriser are insufficient for primary credit, then JAVE be able to accommodate the intuitive lack of knowledge in *The Gettieriser*. But now, by the conditional claim again, the joint contribution of the timeseeker and the timekeeper is insufficient for primary credit and so JAVE predicts lack of knowledge in *The Timekeeper*. Either way, JAVE remains ultimately unsatisfactory.

4.2 Hetherington's Hypothesis of Extended Agents

There is another way of extending epistemology, due to Stephen Hetherington [2012], that may enable us to accommodate the intuition of knowledge in *The Timekeeper*.

According to Hetherington's proposal, the epistemic agent extends beyond the boundaries of persons (i.e., roughly, the biological boundaries of human beings) and may include various aspects of his environment. Hetherington states his 'hypothesis of extended agents' (HEA) in the following passage:

I am suggesting that ... the epistemic agent is the person-plus-more-besides. It can be, for instance, *the person being reliable by consulting a thermometer* – the person-plus-the-thermometer, for a start – when forming a belief about the temperature around her. Even if roughly, the person is thus conceptually distinct from the epistemic agent. *Who, then, is it who knows?* The *person-by-using-the-thermometer* knows; the *person-plus-the-thermometer* knows. That *unity* knows; it is the agent of the knowing.

[Hetherington 2012: 213]

Here is how one might put HEA to use in analysing *The Timekeeper*. Suppose, in keeping with HEA, we allow the epistemic agent

to extend beyond the boundaries of the timeseeker's body. In particular, suppose that the epistemic agent in *The Timekeeper* is not just the timeseeker but the *timeseeker-plus-the-clock-plus-the-timekeeper* (henceforth 'TCT' for short). Suppose, furthermore, that we interpret the agent variables in relevant virtue epistemological and modal conditions on knowledge accordingly. Note that, on this proposal, it is safe to drop the index to methods, as the methods have become part of the agent, as it were.

It is not hard to see that this move will do the trick for modal accounts of knowledge. At the closest worlds at which it is not 8.30 TCT will not believe that it is 8.30. After all, at those worlds the timekeeper alerts the timeseeker of the fact that the clock is inaccurate and the belief is not formed. TCT's belief turns out sensitive after all.

Similarly, TCT avoids false belief whether it's 8.30 across close worlds. Again, at worlds at which it isn't 8.30 the timekeeper will step in and alert the timeseeker to the inaccuracy of the clock. Again the belief that it is 8.30 will not be formed. TCT's belief turns out safe as well.

The situation is slightly less obvious in the case of virtue epistemology. The trick here is that, given that the belief is formed by the extended agent, TCT, the corresponding competence's SI do not include that the clock be functioning properly. If this isn't immediately obvious, notice that, TCT is disposed to perform well cognitively even in SI in which the clock is stopped and so TCT's clock-reading competence may extend to such SI. As a result, TCT can rise to the level of complete competence even when looking at a stopped clock. TCT's belief may qualify as knowledge by the lights of virtue epistemology, too.

Unfortunately, Hetherington's version of HEA account also has some serious drawbacks.

First, suppose we accept this HEA-based analysis of *The Timekeeper*. One question that we have to answer is what goes on in *The Gettieriser*. After all, if we allow the epistemic agent to extend beyond the timeseeker's skin and to include the clock and the gettieriser, it would seem that modal and virtue epistemological accounts that adopt Hetherington's version of HEA will predict, however now incorrectly, that the extended agent's belief is degettiered.

The second problem affects Hetherington's version of HEA more directly. Why are some things agents, whilst other aren't? Why do humans, animals, God and committees but not earthquakes, winds

and natural selection qualify? The overwhelmingly plausible answer is that it has to do with the fact that they possess the capacity for intentional action. By the same token, we would expect that cases of extended agency contribute to the explanation of how the agent so extended possesses the capacity for intentional action. Consider, for instance, Clark and Chalmers's [1998] argument for the extended mind. Suppose they are right and certain mental states are located outside the person's skin. In that case it is also plausible that the agent extends beyond the person's skin. After all, it is plausible that in order to have the capacity for intentional action, one must possess a psychology involving a set of functionally connected mental states in terms of which certain behaviours can be explained as intentional actions. If some of these mental states are located outside the person's skin, parts of their capacity for intentional action will be explained in terms of these external mental states.

Let's now ask whether we have any reason to think that agents are extended in the way Hetherington suggests. What contribution does the hypothesis that, for instance, the agent who knows that it is 19°C at time t and location l is a person-plus-thermometer rather than just a person make to the explanation of the capacity for intentional action? If it could be argued that the mental state involved in the agent's knowledge is held by the extended agent rather than just by the person, we'd have an answer. We could appeal to the above argument from extension of minds to extensions of agency. But now notice just how implausible it is that the relevant mental state is held by the extended agent rather than by the person only. After all, if it were the extended agent rather than the person who holds the mental state, the mental state could not be expected to survive the destruction of the non-personal parts of the extended agent. Clearly, however, it can. One can continue to believe that it is 19°C at t and l even when the thermometer that was involved in the generation of the relevant mental state no longer exists.¹¹ So, the kind of agent extension Hetherington proposes cannot be justified by appeal to the idea that the relevant mental states are held by the extended agent. But then how can it be justified? What other contribution could the external parts of the extended agent make to the explanation of the

¹¹ Notice also how different the situation is in Clark and Chalmers's case of Otto. If Otto's notebook were destroyed, the dispositional beliefs located there would be cease to exist as well. Here it is plausible that it is the extended agent rather than the person who has the mental state.

capacity for intentional action that would be lost if we assumed that the agent were a person only? I cannot see any.

4.3 Extended Methods and TECs

In what follows, I will try to sketch a way of dealing with *The Timekeeper* that, I hope, is more successful than Green's and Hetherington's are. The basic idea is that the case is a case of extended cognition. Crucially, it is the timeseeker's *way of belief-formation* or *method* that extends beyond his skin. More specifically, I want suggest that the timeseeker's method is a complex whole. It includes an on-board part and an external part. While the on-board part is a standard clock-reading process, the external part, which is contributed by the timekeeper, is a *monitoring* process that would alert the timeseeker to inaccuracies of clock readings.

In order to get this proposal off the ground, we will have to answer at least the following three questions:

Question 1. Can we plausibly take the timeseeker's method to be extended in this way?

Question 2. Does the proposal enable modal and virtue epistemological accounts of knowledge to accommodate the intuition of knowledge in *The Timekeeper*?

Question 3. Can we still accommodate the intuition of ignorance in *The Gettieriser*?

Let's first turn to Question 1. Champions of the hypothesis of extended cognition typically accept Clark and Chalmers's 'parity principle' (PP):

If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process.

[Clark and Chalmers 1998: 8]

As far as PP is concerned, there is no problem for the present strategy. If the timeseeker had (and used) an internal monitoring process of the sort now contributed by the timekeeper, we would have no qualms about accepting it as part of his way of belief formation or method. If the timeseeker had been fitted with an internal time measuring

device that would alert him to inaccuracies of external-world clocks, we would have no hesitation in accepting it as part of his way of forming beliefs about the time.¹²

A difficulty may yet arise from another requirement on extended cognition that Clark and Chalmers countenance. They acknowledge that coupled systems (i.e. cognitive systems in which some of the processing is done externally) require the *reliable coupling* of the internal and the relevant external components. Moreover, the relevant sense of reliability they appear to have in mind is unpacked in terms of how easily the coupled system may have been decoupled [Clark and Chalmers 1998: 10-11]. It is not hard to see that this requirement poses a real difficulty for the present proposal. After all, the coupling between the timeseeker's clock-reading process and the timekeeper's monitoring process is highly fragile and the system is very easily decoupled (in fact it will likely be decoupled before the timeseeker uses his clock-reading process again).

Fortunately, at least Clark has backtracked from the reliability requirement. In a more recent paper, Wilson and Clark argue that there may be cases of "extended systems that involve temporary, transient forms of cognitive augmentation", which they call "Transient Extended Cognitive Systems" or "TECSs" [Wilson and Clark 2009: 65]. Wilson and Clark are very clear that, while there will be TECSs involving reliable couplings with the extended parts, there may also be short-lived TECSs, featuring only one-off couplings. Now, if we allow for such short-lived TECSs, the prospects for the present proposal become brighter again.

What does matter to whether we are dealing with an extended system is that "it needs to achieve functional integrity when operating together with the rest of some cognitive system that serves the kinds of purposes that that cognitive system has served". [Wilson and Clark 2009: 63] Notice that this is very plausibly achieved here. The purpose of the timeseeker's clock-reading method is to reliably

¹² It may be worth noting that, use of such monitoring processes is no exception. I look at a pendulum clock, see that it reads 2.30 and yet do not form the corresponding belief: I also see that the pendulum is not moving. I use a monitoring process that alerts me to the fact that the clock is likely inaccurate. Notice that beliefs about the time I form by looking at pendulum clocks are regularly formed via a method including this monitoring process even when it does not become active. For that reason, when I form a belief that it is 2.30 on the basis of a reading of a pendulum clock with a moving pendulum, my belief is formed via a method that includes a monitoring process of the same general kind as the one contributed by the timekeeper in *The Timekeeper*.

lead her to true beliefs (or perhaps knowledge) about the time. The contribution that a monitoring process makes to the method is to enhance the reliability by preventing the formation of beliefs (or by retracting existing beliefs) in conditions in which the first-order process is unreliable. In *The Timekeeper*, the external monitoring process achieves this just as effectively and smoothly as their familiar on-board cousins. Functional integrity is plausibly achieved here.

Since the timeseeker's method satisfies PP and since it achieves the relevant functional integrity, we can, with Wilson and Clark, view it as a TECS including an on-board clock-reading process and an external monitoring process. The answer to Question 1, then, is positive.¹³

Let's move on to Question 2. Given that the timeseeker's method or way of belief-formation is such a TECS, can we accommodate the intuition of knowledge?

It is not hard to see that, for modal accounts of knowledge, the answer is 'yes'. At the closest worlds at which it is not 8.30 and the timeseeker uses this TECS to arrive at a belief whether it is 8.30, he will not believe via the extended method that it is 8.30. After all, at those worlds the timekeeper will have alerted her of the fact that the clock is stopped. Hence, if the timeseeker's method is extended in the way envisaged, her belief turns out sensitive after all.

Similarly, at close worlds at which the timeseeker uses the TECS, she does not arrive at a false belief whether it's 8.30. Again, at worlds

¹³ It may be worth mentioning that, even among champions of the hypothesis of extended cognition, the question whether transient forms of cognitive extension are possible remains a matter of controversy. (Thanks to Cameron Boult for pointing this out to me.) For the purposes of my argument, I will assume that this controversy can be resolved in a favourable way and will leave it to my colleagues in the philosophy of mind to do so. That said, there is some reason to think that the transient kinds of coupling I need here are possible when it comes to the ways of performing at issue in abilities. An archer may have the ability to hit the target in certain IN and SI only with very rare equipment, say, a type of bow of which only one copy exists. It may also be that the coupling between archer and the bow is highly unreliable: The archer has been granted use of it by the owner on a whim and for a very short period of time only. Even so, when the archer shoots with the bow in IN and SI, he may rise to the level of complete competence and, as a result, may attain successes that manifest competence. In the case of the ways of performing at issue in abilities, then, unreliable coupling with crucial equipment is compatible with complete competence. If so, one might think that the same holds for ways of belief formation. So long as the coupling is achieved, whether reliably or not, the agent may rise to the level of complete competence, in which case, by the lights of VE at least, the beliefs he forms may qualify as knowledge (see also below).

at which it isn't 8.30 the timekeeper will step in and alert her to the fact that the clock is stopped. Hence, if the timeseeker's method is the TECS, her belief turns out safe.

To see how VE can accommodate the intuition of knowledge, note that given that the way of belief formation is the relevant TECS, the corresponding competence's SI do not include that the clock be functioning properly. After all, using the TECS disposes the agent to perform well cognitively even in SI in which the clock is stopped. As a result, the timeseeker can rise to the level of complete competence even when he is looking at a stopped clock. Her belief may qualify as knowledge by the lights of virtue epistemology as well. We thus also have a positive answer to Question 2.

This leaves Question 3, i.e. the question whether we can also accommodate the intuition of ignorance in *The Gettieriser*. Notice that if the timeseeker's method is exhausted by his on-board processes, we will be able to secure a positive answer. After all, in that case, *The Gettieriser* is in all relevant respects analogous to *The Original*. And we have already seen this case does not pose a significant difficulty for modal or virtue epistemological accounts of knowledge. On the other hand, a difficulty looms if the timeseeker's method extends beyond his skin and includes the doings of the gettieriser in much the same way as, by the above analysis of *The Timekeeper*, her method extends to include the monitoring process contributed by the timekeeper. After all, in that case, it would seem as though the beliefs will satisfy the anti-Gettier conditions of our modal and virtue epistemological accounts. The accounts predict, mistakenly, that the timeseeker's belief is not gettiered.

So let's ask whether the timeseeker's method in *The Gettieriser* can plausibly be thought to extend beyond her skin in this way. Consider, once more PP: If, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is part of the cognitive process. Do the gettieriser's doings satisfy PP? One might think it does not. After all, how could the setting of some clock occur in the timeseeker's head? Well, suppose that the timeseeker had a brain-implant that happens to interfere with the clock in much the same way as the gettieriser does in the problematic case. But even so, the process is not occurring in the timeseeker's head. After all, the clock remains an external object. So suppose that the stopped clock the timeseeker consults is

an internal (e.g. mental) clock. Unbeknownst to the timeseeker, the internal clock is stopped and yet displays the time accurately because another cognitive (sub)system happens to set it to the right time just before the timeseeker consults it.

This is the best attempt at offering an internal analogue to *The Gettieriser* that I can think of. At the same time, I must confess that I am not quite sure what to make of it. The case is so bizarre that I have little inclination to pass any verdict of philosophical substance whatsoever about it, and certainly not *without hesitation*. If I am not the only one feels this way about the case, the prospects for a positive answer to Question 3 are starting to looking up. After all, in that case, PP does not enjoin us to view the timeseeker's method as extending to the gettieriser. If it doesn't, *The Gettieriser* is structurally analogous to *The Original* and so does not pose any special difficulties for the present proposal.¹⁴

Conclusion

If the argument I have developed in this paper goes through, we have some reason to opt for an extended rather than a traditional (non-extended) epistemology. There are cases, notably *The Timekeeper*, that pose a problem for even our most promising traditional accounts of knowledge, but can be dealt with nicely by (at least a certain type of)

¹⁴ Of course, PP offers only a sufficient condition for extension, not a necessary one, and so we have no conclusive evidence that the timeseeker's method does not extend to the gettieriser. At the same time, given that PP is all we have to work with, we at least have no reason to think that the timeseeker's method does extend to the gettieriser. So, we can leave it to critics of the present account to produce such a reason. In addition, there is reason to think that champions of the extended cognition hypothesis will do well to try and avoid allowing extension of the timeseeker's method to the gettieriser. After all, for any version of the extended cognition hypothesis on which the timeseeker's method does so extend, even modal and virtue epistemological accounts of knowledge will predict, incorrectly, that the timeseeker's belief is not gettiered. That means that even our most promising accounts of the anti-Gettier condition are bound to fail. The envisaged versions of the extended cognition hypothesis would turn out to have highly undesirable epistemological consequences. As epistemologists, then, we are in the fortunate position of leaving it to our critics to produce reason to believe that the timeseeker's method extends to the gettieriser in the potentially problematic way and to our colleagues in the philosophy of mind to remedy the defect should our critics succeed. Of course, if all of this came to pass, we could no longer be assured that the proposed account of *The Timekeeper* will continue to work in the way envisaged. Even so, there is not much we can do about this now. We will have to wait for our critics to make their move.

extended epistemology. Notice that the case also serves to provide epistemological evidence in support of the extended cognition hypothesis. Given that the timeseeker has knowledge and the only way to accommodate this datum is by appeal to cognitive extension, there is epistemological evidence that the hypothesis of extended cognition is true. On the other hand, given that the timeseeker's method extends to the timekeeper, there is also reason to believe that the timeseeker does indeed acquire knowledge in this case. After all, in that case, her belief not only satisfies a number of our most promising degettierisation conditions on knowledge, but it is also acquired via a method the internal counterpart of which would uncontroversially deliver knowledge. Once we accept the hypothesis of extended cognition, it would seem to be mere prejudice not to regard this case as a case of knowledge. To the extent that the hypothesis of extended cognition is independently plausible, then, we have evidence from the philosophy of mind that the timeseeker does really know in this case. In this way, the intuition that, in *The Timekeeper*, the timeseeker knows and the extended cognition hypothesis constitute mutually reinforcing strands of evidence.

Once we opt for an extended epistemology *The Timekeeper* becomes tractable. In fact, I have argued that there are a number of extended epistemologies that can get this case right. However, these views encounter difficulties elsewhere. The task that extended epistemologies face is to avoid extensions to occur too easily. More specifically, they must avoid them in certain types of Gettier case such as *The Gettieriser*. I have shown that two existing versions of extended epistemology, to wit, Green's joint abilities and Hetherington's extended agents account, fail on this count. Fortunately, there is reason to think that there is a workable kind of extended epistemology after all. I have sketched one such view on which what extends is the agent's method and I have argued that it can retain the benefits an extended epistemology promises, whilst avoiding at least some of the pitfalls this type of view encounters.¹⁵

¹⁵ *Acknowledgements.* Many thanks to Harmen Ghijsen, Sandy Goldberg, Jesper Kallestrup, Duncan Pritchard and the audience of the *Second Workshop on Extended Epistemology* at the University of Edinburgh for helpful discussions of this paper. Special thanks to Orestis Palermos and Cameron Boulton for their excellent comments on an earlier draft of the paper.

References

- Becker, K. 2007. *Epistemology Modalized*. Routledge, New York.
- Becker, K. 2009. Margins for error and sensitivity: what Nozick might have said. *Acta Analytica*, 24:17-31.
- Black, T. 2008. Defending a sensitive Neo-Mooreanism. In Hendricks, V. and Pritchard, D., editors, *New Waves in Epistemology*. Palgrave Macmillan, Basingstoke.
- Black, T. and Murphy, P. 2007. In defense of sensitivity. *Synthese*, 154:53-71.
- Clark, A. and Chalmers, D. 1998. The extended mind. *Analysis*, 58:7-19.
- Dretske, F. 1971. Conclusive reasons. *Australasian Journal of Philosophy*, 49:1-22.
- Dretske, F. 1981. *Knowledge and the Flow of Information*. MIT Press, Cambridge/MA.
- Gettier, E. 1963. Is justified true belief knowledge? *Analysis*, 23:121-23.
- Goldberg, S. 2010. *Relying on Others: An Essay in Epistemology*. Oxford University Press, New York.
- Greco, J. 2003. Knowledge as credit for true belief. In DePaul, M. and Zagzebski, L., editors, *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford University Press, Oxford.
- Greco, J. 2010. *Achieving Knowledge*. Cambridge University Press, Cambridge.
- Greco, J. 2012. A (different) virtue epistemology. *Philosophy and Phenomenological Research*, 85:1-26.
- Green, A. 2012. Extending the credit theory of knowledge. *Philosophical Explorations*, 15:121-32.

- Hetherington, S. 1999. Knowing failably. *The Journal of Philosophy*, 96:565-87.
- Hetherington, S. 2012. The extended knower. *Philosophical Explorations*, 15:207-18.
- Kelp, C. 2011. In defence of virtue epistemology. *Synthese*, 179:409-33.
- Kelp 2013a. Extended cognition and robust virtue epistemology. *Erkenntnis*, 78:245-52.
- Kelp, C. 2013b. Extended cognition and robust virtue epistemology: Response to Vaesen. Forthcoming in *Erkenntnis*.
- Kelp, C. 2013c. Knowledge: The safe-apt view. *Australasian Journal of Philosophy*, 91:265-78.
- Kelp, C. 2013d. Knowledge, understanding and virtue. Forthcoming in Fairweather, A., editor, *Virtue Scientia. Virtue Epistemology and Philosophy of Science*. Springer, Dordrecht.
- Lackey, J. 2007. Why we don't deserve credit for everything we know. *Synthese*, 158:345-61.
- Lackey, J. 2009. Knowledge and credit. *Philosophical Studies*, 142:27-42.
- Nozick, R. 1981. *Philosophical Explanations*. Oxford University Press, Oxford.
- Pritchard, D. 2005. *Epistemic Luck*. Oxford University Press, Oxford.
- Pritchard, D. 2007. Anti-luck epistemology. *Synthese*, 158:277-297.
- Pritchard, D. 2008. Greco on knowledge: Virtues, contexts, achievements. *The Philosophical Quarterly*, 58:437-47.
- Pritchard, D. 2010. Cognitive ability and the extended cognition thesis. *Synthese*, 175:133-51.

- Pritchard, D. 2012. Anti-luck virtue epistemology. *The Journal of Philosophy*, 109:247-79.
- Pritchard, D., Millar, A., and Haddock, A. 2010. *The Nature and Value of Knowledge*. Oxford University Press, Oxford.
- Riggs, W. 2002. Reliability and the value of knowledge. *Philosophy and Phenomenological Research*, 64:79-96.
- Riggs, W. 2009a. Luck, knowledge, and control. In Haddock, A., Millar, A., and Pritchard, D., editors, *Epistemic Value*. Oxford University Press, Oxford.
- Riggs, W. 2009b. Two problems of easy credit. *Synthese*, 169:201-216.
- Sainsbury, M. 1997. Easy possibilities. *Philosophy and Phenomenological Research*, 57:907-19.
- Sosa, E. 1999. How to defeat opposition to Moore. In Tomberlin, J., editor, *Philosophical Perspectives*. Blackwell, Oxford.
- Sosa, E. 2007. *A Virtue Epistemology. Apt Belief and Reflective Knowledge*, volume 1. Oxford University Press, Oxford.
- Sosa, E. 2010. How competence matters in epistemology. *Philosophical Perspectives*, 24:465-475.
- Sosa, E. 2011. *Knowing Full-Well*. Princeton University Press, Princeton/NJ.
- Vaesen, K. 2011. Knowledge without credit, exhibit 4: extended cognition. *Synthese*, 181:515-29.
- Vaesen, K. 2013. Critical discussion: virtue epistemology and extended cognition: a reply to Kelp and Greco. *Erkenntnis*, 78:963-70.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford University Press, Oxford.

Wilson, R. and Clark, A. 2009. How to situate cognition: Letting nature take its course. In Aydede, M. and Robbis, P., editors, *The Cambridge Handbook of Situated Cognition*. Cambridge University Press, Cambridge.

Zagzebski, L. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press, Cambridge.