Title

Epistemically useful false beliefs

Permalink

https://escholarship.org/uc/item/7qt4t3qc

Journal

Philosophical Explorations, 20(sup1)

ISSN

1386-9795

Author

Pritchard, Duncan

Publication Date

2017-03-24

DOI

10.1080/13869795.2017.1287291

Copyright Information

Peer reviewed

EPISTEMICALLY USEFUL FALSE BELIEFS

Duncan Pritchard

School of Philosophy, Psychology *&* Language Sciences

University of Edinburgh

duncan.pritchard@ed.ac.uk

ABSTRACT. Our interest is in the possibility of there being a philosophically interesting set of useful false beliefs where the utility in question is specifically epistemic. As we will see, it is hard to delineate plausible candidates in this regard, though several are promising at first blush. We begin with the kind of strictly false claims that are said to be often involved in good scientific practice, such as through the use of idealisations and fictions. The problem is that it is difficult to see that there would be any epistemic utility in believing such claims, as opposed, say, to merely accepting them. Next we turn to the challenge posed by *epistemic situationism*, which when embedded within a plausible form of virtue epistemology appears to show that sometimes purely situational factors can play a significant explanatory role in one's cognitive success. But again it is hard to see how the role that these epistemically beneficial situational factors contribute can be cashed out in terms of epistemically useful false beliefs on the part of the subject. Finally, we turn to the Wittgensteinian conception of *hinge commitments*, commitments that are held to be epistemically useful even if false. While the epistemic utility of these commitments is defended, it is argued that one cannot make sense of these commitments in terms of belief. Support is thus canvassed, albeit in a piecemeal fashion, for the thesis that the prospects for there being philosophically interesting cases of epistemically useful false belief are poor.

BIO: Duncan Pritchard is Professor of Philosophy at the University of Edinburgh, and Director of the *Eidyn* research centre. His monographs include *Epistemic Luck* (Oxford UP, 2005), *The Nature and Value of Knowledge* (co-authored, Oxford UP, 2010), *Epistemological Disjunctivism* (Oxford UP, 2012), and *Epistemic Angst* (Princeton UP, 2015).

1. EPISTEMICALLY USEFUL FALSE BELIEFS

Our interest is in the possibility of there being epistemically useful false beliefs. This means that four conditions need to be met:

(i)   the subject genuinely *believes* the target proposition;[1]
(ii)   the belief is *false*;
(iii)   the belief is *useful*;
(iv)   the utility is specifically *epistemic*.

A few comments are in order regarding these conditions.

On condition (i), it is important to emphasise that we are interested in the notion of belief that is of concern to epistemologists—*viz*., that propositional attitude which is a constituent part of knowledge. This point is important since the notion of belief is used in a multitude of ways, some of them more permissive than others.[2] But the notion of belief that is operative in epistemology is specifically concerned with a commitment to the truth of the target proposition (i.e., such that one believes that *p is true*). This has some important implications for our purposes. For example, belief in this sense is incompatible with an attitude of agnosticism regarding the truth of *p*. It also means that this is a propositional attitude which has some basic conceptual connections to truth and rationality. In particular, if one recognises that one has no rational basis at all for regarding *p* as true, then one's commitment to that proposition cannot be one of belief in this sense, but is rather a different propositional attitude (e.g., one of wishful thinking that *p*).[3] Henceforth, when I discuss the notion of belief without qualification, it will be this specific conception of belief that I will have in mind.

Notice how condition (i), as we are understanding it, effectively rule-outs certain potential cases of epistemically useful true belief. For example, it cannot be the case that there are epistemically useful false beliefs where the agent is aware that the belief is false, since in that case the propositional attitude wouldn't quality as a belief in the sense that we have in mind. For the same reason, there cannot be cases of epistemically useful false beliefs where the agent recognises that she has no rational basis for the truth of the target proposition, or where the agent is agnostic about the truth of the target proposition (regardless of whether this is because she recognises that she has no rational basis for its truth).

Next, consider conditions (iii) and (iv). There are various potential ways in which false beliefs can be useful, particularly if the utility in question is merely practical. It may, for example, be a sign of good mental health that one tends to overestimate one's abilities (or how good-looking one is, and so on). There is at least a *prima facie* case in the relevant empirical

literature that false beliefs of this kind will tend to enable one to be a happier person.[4] But our interest here is not in useful false beliefs of this general kind, but in the more specific question of whether there can be *epistemically* useful false beliefs. This claim is far more controversial, and for good reason. After all, the mark of the epistemic is typically held to be that it is truth-conducive, so how could a false belief have any epistemic utility?[5]

We need to introduce a further condition into our analysis in order to bring this issue into sharper relief. After all, that a false belief might have some short-term epistemic benefit is obviously true, and doesn't in any way run counter to the idea that the mark of the epistemic is that it is truth-conducive. This is because this latter idea is obviously best understood over the longer-term, rather than the short-term.

Suppose, for example, that one makes some errors in one's calculations that, as it happens, enables one to form a true arithmetical belief that one wouldn't have gained otherwise (perhaps the errors cancelled each other out, but that absent the errors this arithmetical approach to the problem would have delivered the wrong verdict). This would thus be a case of short-term epistemically useful false belief. But one would naturally expect the errors in play here to lead to false beliefs over the longer term, and hence that there would be epistemic disvalue present which cancels out the short-term epistemic gain. That's why, from an epistemic perspective, we care to eliminate such errors, since they are not conducive to the truth. Our concern is thus not with the trivial cases of epistemically useful false beliefs where the epistemic utility is short-term, but rather in the possibility of more interesting cases where there is a long-term epistemic benefit to having false beliefs. We can thus add the following fifth condition concerning what we are looking for in an example of epistemically useful false belief:

(v)   the epistemic utility is assessed over the *long-term*.

There is a related point we need to make in this regard, which is that we are specifically interested in false beliefs that are epistemically useful in a systematic fashion that arises out of the false belief itself, as opposed to the epistemic utility merely being, for example, a matter of happenstance. After all, it is very easy to engineer cases of epistemically useful false belief, even over the long term, if don't add this stipulation. Imagine, for example, that one is a kind of epistemic 'Mr. Magoo', whereby everything happens to work out from an epistemic point of view, even though one has many false beliefs, but where this is all down to luck.[6] Here there is a sense in which one's false beliefs might well be epistemically useful, even over the long-term, but it isn't a particularly interesting case, as the epistemic good fortune has only an incidental connection to one's formation of the false belief.

Indeed, even if the epistemic utility occurs in a systematic fashion, if it doesn't arise out of the false belief in question then the case is equally uninteresting for our purposes. Imagine, for example, that one has a helpful demon who is devoted to ensuring that one's false beliefs always generate good epistemic consequences. Here the epistemic utility is systematic, rather than being due to good fortune, but it has nothing essentially to do with one's formation of the false belief. It is thus a trivial case of epistemically valuable false belief, in that the epistemic value in question has nothing essentially to do with the target false belief. We thus get our sixth condition concerning what we are looking for in an example of epistemically useful false belief:

(vi)   the epistemic utility *systematically* arises out of the false belief itself.

A final point is in order, regarding how we are to assess epistemic utility (or disutility). It is tempting to do so by simply doing a reckoning of a subject's true and false beliefs. So, for example, if having a false belief that p leads the subject overall to have more true beliefs—or at least a better ratio of true beliefs over false beliefs—than she would have had otherwise, then the false belief that p is epistemically valuable. While I admit that such an approach might be initially appealing in virtue of its simplicity, I think it would be a disastrous way to proceed. This is because a moment's reflection reveals that we cannot possibly 'weigh' epistemic value in such a crude manner.

Suppose one is faced with two doors, where behind each can be found a true proposition about which one is currently unaware. But suppose further that while one of the true propositions concerns something entirely trivial, such as how many hairs one has on one's head, the other true proposition concerns something of great import, such as the key to unlocking a unified theory that makes general relativity and quantum field theory compatible with one another. Clearly, one would prefer to open the door with the substantial truth behind it rather than the trivial one. Crucially, however, the reason for this is precisely that one cares about the truth, in that one's grip on the nature of reality is substantially improved by learning the deep scientific truth in a way that it isn't improved when one learns the trivial truth. That is, while both propositions are true, there is a perfectly legitimate sense in which there is *more* truth contained in the substantial truth than in the trivial truth.[7]

Notice that my point here is entirely a negative one—*viz*., that one cannot simply assess epistemic value (or disvalue) by counting the true (or false) propositions that a subject believes. In particular, I'm not proposing a positive account of how one should instead assess epistemic value. My view is that this is a deep metaphysical problem that has not yet been properly resolved, and it is certainly an issue which would take us too far afield to explore here.[8] This

negative point does have implications for how we should understand conditions (iv) and (v), however, in that we should be wary of being too quick to infer from the fact that a false belief leads to true beliefs over the long-term that it is thereby an example of long-term epistemic utility, as the former does not guarantee the latter.[9]

I intend to approach the question of whether there can be epistemically useful false beliefs in a piecemeal fashion, by considering a selection of plausible candidates. These include the use of idealisations and fictions in science, the kinds of cases cited by proponents of epistemic situationism, and the notion of a hinge commitment as it appears in Wittgenstein's later work. As we will see, each of these cases ultimately fails to be an instance of epistemically useful false belief, either because it is not epistemically useful in the required fashion, or because it fails to be a genuine belief in the sense that we are interested in. Naturally, such a piecemeal approach to the issue does not settle the matter of whether there can be such beliefs, but it will at least help us to articulate why such a notion is inherently problematic, and thereby establish a *prima face* case for thinking that there cannot be beliefs of this kind.

2. SCIENTIFIC IDEALISATIONS AND FICTIONS

We will begin with Catherine Elgin's (1996; 2004; 2009) seminal work on idealisations in scientific understanding. Her claim is that much of good scientific practice involves the use of models and idealisations that are, strictly speaking, false. Think, for example, of the ideal gas law, which is a very useful way of understanding the behaviour of gases but is also, *qua* an idealisation, not strictly true. This case is far from isolated, as Elgin's work is full of examples that demonstrate that scientific progress often appeals to claims which are not strictly true. The nub of the matter for our purposes is that a well-conducted scientific inquiry—particularly one that has demonstrably furthered scientific progress, and hence enhanced our scientific understanding of the world around us—is surely an excellent guide to epistemic utility. Hence if false scientific beliefs can often be beneficial for scientific inquiries, then it follows that false scientific beliefs can be epistemically useful.[10]

One problem with this account of epistemically useful false belief, however, is that it is far from obvious that the subjects concerned really believe the falsehood in question. After all, don't the scientists know full well that the ideal gas law is an idealisation? If so, then it seems that Elgin is using the notion of belief in a much broad sense to that which we delineated above, such that it can apply to a propositional attitude which is consistent with the subject

concerned actually disavowing the truth of the target proposition. In particular, it seems that what Elgin really has in mind in such cases is not a believing at all but rather something weaker, like an *acceptance*. Unlike belief (as we are understanding that notion anyway), one can accept that *p* while being agnostic about whether *p* is true. Often in scientific inquiry the right propositional attitude to a theoretical claim is acceptance rather than belief—e.g., where the best science available suggest *p*, but where the evidential case for *p* is highly defeasible, such that one recognises that much more scientific investigation is required to establish *p*.[11]

Acceptance that *p* is also surely the relevant propositional attitude to take when one recognises that while *p* is false, there is a clear epistemic utility to acting as if it were true, in that it usefully approximates the truth. Isn't that just the propositional attitude manifested by scientists who employ the ideal gas law in their scientific reasoning? Indeed, one might think that, if anything, the propositional attitude in play here is even weaker than acceptance that *p*, but is rather something like a *willingness to suppose* that *p*. Accepting that *p* often correlates with the view that *p* is *likely* to be true, after all, and yet a willingness to suppose that *p* is compatible with even the idea that *p* is strictly speaking false. In any case, the crucial point is that acceptance that *p* is not a form of believing given the point we made earlier about belief that *p* excluding agnosticism that *p*. Hence it follows that Elgin hasn't given us a case of epistemically useful false belief.

Of course, one might argue, in defence of Elgin, that there could be cases of practicing scientists who don't realize that the ideal gas law is an idealisation, and hence don't realize that it is strictly speaking false. These scientists may actually believe a falsehood, where there is an epistemic utility in them having this belief in virtue of the role that it plays in promoting their scientific inquiry. Is this a legitimate case of epistemically useful false belief?

I don't think so. To begin with, notice that the epistemic utility of this belief, now that we grant that it a genuine false belief, starts to look somewhat dubious. After all, as we were previously thinking of the case, it was the fact that this was a roughly accurate—but recognised to be strictly speaking false—proposition that was guiding the scientific inquiry. There is no error in play here under this description, as the scientist is fully aware that she is working with an idealisation. Now, however, we have something very different, in that we are supposing that the scientist is actually in error regarding the relevant facts. Would this really *promote* their scientific inquiry? In the case of the ideal gas law, my understanding is that whether the scientist believes the fiction in play, or is merely willing to accept it (while recognising that it is a useful fiction), will make no practical difference to their scientific inquiries. But in that case, to what extent is this an *epistemically* useful false belief, given that no positive epistemic utility is

generated by one having it? Moreover, although this example might concern a useful fiction where false belief rather than mere acceptance makes no practical difference, there might be other cases where actively falsely believing the fiction actually has a deleterious effect on one's scientific inquiry, perhaps by closing off theoretical options that should have been kept open (and would have been, had the subject been aware that the useful fiction is just that). In this case, the false belief in question has an epistemic *dis*value, and hence we no longer even have a case of epistemically useful false belief anyway.

One might respond to this critical line of reasoning by saying that it is unfair to contrast false belief with mere acceptance (in the latter case in conjunction with a recognition that to believe would be to believe falsely). Instead, why not contrast having the false belief and proceeding accordingly with having no belief at all, and hence not employing the scientifically useful fiction? Wouldn't it be of epistemic value to have the false belief if the alternative is to have no belief at all, and hence not to proceed down a promising scientific avenue?

My view is that what one gains in such a trade-off is at most short-term epistemic value at the expense of a longer-term epistemic disvalue (one that will, at least ordinarily, cancel out the short-term epistemic value), and yet, as we noted above, it is the putative long-term epistemic value of a false belief that we are interested in. I think we can see why this is so by considering a parallel phenomenon in educational theory. Educational theorists often talk of the role of educators as providing 'scaffolded' learning environments, where these scaffolds are normally understood in broadly epistemic terms.[12] That is, the educator is assisting the pupil in their epistemic endeavours, sometimes in ways that the student is unaware.[13] This needn't involve subterfuge, of course, as often the educator will be quite explicit about how they are helping the student. But it might, and when it does it may well result in the pupil forming false beliefs, at least over the short-term.

For example, when teaching a subject matter that is highly complex, such as evolutionary theory, it may be preferable to begin with some rough statements of core elements of the theory in order for the student to gain a handle on what the theory proposes. It is quite possible that a student might in the process acquire some false beliefs as a result—e.g., that human beings literally are monkeys, say—and it may be a sound educational decision on the part of the educator to let such a false belief slide over the short term. After all, one wants to exploit their enthusiasm to learn more about the subject matter, and this might be best-served by focussing on what they have got right as opposed to the few things that they have got wrong, rather than delving too quickly into more complex parts of the view.[14]

Still, no conscientious educator would be content to leave the matter there, and for good reason. Even if turning a blind eye to a false belief might be epistemically legitimate in the short term—because it is epistemically useful, over the short term—it would be entirely inimical to the epistemic goals of the educational enterprise to let such a false belief go uncorrected over the longer term. Moreover, this is not a mere fetishism for the truth on the part of the educator, but rather reflects a recognition of how problematic false belief can be if allowed to gain roots. In particular, while the false belief might have short-term educational advantages, over the longer term it will undermine the student's capacity to truly get to grips with the subject matter in question.[15]

I think that the educator's stance towards allowing, over the short-term, some false beliefs to be acquired is similar to how we should view the putative epistemic advantages of false beliefs in the scientific case. In particular, this is an epistemic benefit that only makes sense over the short-term, in that error will tend to lead to further error over the longer term. Thus, just as we might be sanguine about false belief in a student as part of a developmental stage that enables her to reach higher educational goals, so we might well be sanguine about a scientist who has false beliefs over the shorter term if this is the only way that she might make any scientific progress. But such false belief is no rational resting-place over the long term, either for the developmental student or the 'developmental' scientist, and the reason for this is that false beliefs tend to be epistemically disvaluable over the longer term.

Consider, for example, the case of a 'developmental' scientist who employs Newtonian mechanics while oblivious to the fact that the principles in play in this regard are not strictly speaking correct, in that, for example, they fail to deliver the right results when applied at the subatomic level. Even so, our developmental scientist may only ever employ these principles with regard to medium-sized objects where they will reliably deliver the right verdicts. Over the short term, then, there will thus be an epistemic advantage to having false beliefs in these principles when compared with the alternative scenario where the subject lacks any beliefs at all about Newtonian mechanics (and where the subject doesn't accept the relevant principles, or is otherwise willing to suppose to that they are true).

But over the long term—which as we noted above, is the time-frame that concerns us—the epistemic benefits of these false beliefs will recede, and the epistemic costs will mount. After all, insofar as our subject retains a conviction in the truth of these principles then this will close-off important theoretical options and thereby stymie scientific progress. Moreover, we also need to keep in mind our point from the previous section that we should be wary of measuring epistemic utility simply in terms of the production of true beliefs. Employing these

Newtonian principles might well lead to lots of true beliefs, but these will all be true beliefs of a similar kind, and predicated upon an ultimately inaccurate conception of the nature of the physical world. It is far better, from an epistemic point of view, to have a more accurate conception of the world around one, such that one is able to form a wider range of true beliefs in such a fashion that is not predicated upon an inaccurate conception. As we might put the point, someone who understands the limitations of Newtonian mechanics has a far greater epistemic grip on the nature of the physical world than one who lacks this understanding. This is a kind of epistemic value that is not captured by simply counting the true beliefs that these agents go on to form as a result of their very different conceptions of the world around them. In particular, even if a lack of such an understanding—and the possession of the relevant false beliefs—leads to a greater number of true beliefs, this doesn't ensure that these false beliefs are epistemically valuable.

Furthermore, we also need to keep in mind the additional point that we are interested in epistemic value that systematically arises out of the false belief in question, rather than being, for example, due to mere happenstance. Even if, as it happens, one's use of these Newtonian principles never leads one astray, this is only because of one's limited range of application of these principles. Had one applied these principles to a broader range of propositions, for example, then one would have been led astray, and hence have formed false beliefs as a consequence. The crux of the matter is that insofar as there is epistemic utility being generated here, it is not a form of utility that is systematically arising out of the false beliefs in question, but rather turns on the good fortune that one only employs these false beliefs within a restricted range where they will not lead to error.

Putting all these points together, the prospects of offering an interesting form of epistemically useful false belief from instances of scientific idealisation and fictions do not look favourable.

3. EPISTEMIC SITUATIONISM

A more promising case of epistemically useful false belief may be provided by *epistemic situationism*. As a general view, situationism emphasises results from empirical psychology that appear to show that apparently incidental—or *situational*—factors can play a hitherto unrecognised significant role in one's agency, as when certain smells can have a dramatic impact on behaviour (e.g., making one more generous to beggars). In ethics, these empirical

results have been used to claim that virtue theory is untenable, in that it is often not broad character traits like virtues that guide moral behaviour but rather merely these situational factors.[16] But the view also has potential application to epistemology, and hence we get epistemic situationism. After all, the dominant brand of contemporary epistemology—*virtue epistemology*—incorporates the idea that knowledge and other positive epistemic standings are the result (in some significant way) of cognitive agency, and yet the situationist empirical literature seems to suggest that many of these cognitive successes are in fact due to situational factors rather than cognitive agency.[17]

As I've argued elsewhere—see, for example, Pritchard (2014*a*) and Carter & Pritchard (*forthcoming*)—the challenge posed to virtue epistemology by epistemic situationism only really applies (at least with the intended force) to one version of the view, a version which is independently implausible. According to *robust virtue epistemology*—as defended, in various guises, by Ernest Sosa (1991; 2007; 2009; 2016), Linda Zagzebski (1996; 1999), and John Greco (2009)—knowledge is to be understood as cognitive success (i.e., true belief) which is the manifestation of cognitive ability (i.e., epistemic virtue, broadly understood). Crucially, on this view, there is meant to be no need for any further epistemic conditions on knowledge, such as an anti-luck condition. This means that the virtue requirement on knowledge needs to be fairly austere, in order to deal with standard cases of Gettier-style epistemic luck. Such cases, after all, usually involve a fairly high degree of cognitive ability on the part of the subject, and yet the subject nonetheless fails to acquire knowledge, even though she is cognitively successful, on account of the knowledge-undermining epistemic luck in play. To exclude such cases with the virtue condition on knowledge, the proponent of robust virtue epistemology needs to demand that the subject's manifestation of cognitive ability should be the overarching element in the causal explanation of her cognitive success, since that is something that isn't exhibited in Gettier-style cases (on account of the fact that it is the luck involved that is carrying most of the explanatory burden).

If the empirical literature cited by epistemic situationism is sound, then robust virtue epistemology would seem to entail that we know an awful lot less than we ordinarily suppose. In particular, this literature appears to suggest that often our cognitive abilities are not playing as much of a role in accounting for our cognitive success as we imagine, with situational factors also playing a significant part in producing that cognitive success. That's not a count against robust virtue epistemology in itself, of course, since perhaps it is simply true that we lack a lot of the knowledge that we ordinarily tend to attribute to ourselves. Once we understand the shortcomings of robust virtue epistemology, however, then it becomes clear that virtue

epistemology (of a more modest form) can embrace epistemic situationism without this leading to a mitigated form of scepticism about the extent of our knowledge.

The first point we need to make in this regard is to be clear on the nature of the empirical claims offered by epistemic situationism. What this literature does clearly show is that situational factors can often play a much larger role in our cognitive successes than we hitherto imagined, which of course has the consequence that our cognitive abilities are playing less of a role than we hitherto imagined. What the literature doesn't show, however, is that our cognitive successes are often exclusively down to situational factors rather than our cognitive abilities. The former claim is in tension with robust virtue epistemology, at least if we want to avoid placing restrictions on the scope of our knowledge. But it is not in tension with a *modest virtue epistemology* that argues only for the necessity of a subject meeting a weak virtue condition on knowledge (i.e., where the virtue condition is not the only condition on knowledge).[18]

This point is important because there are independent reasons for preferring modest virtue epistemology to its robust counterpart. In particular, only modest virtue epistemology is able to account for what I term the *epistemic dependence* of knowledge. This concerns the way in which whether one's cognitive success counts as knowledge can be significantly dependent on factors outwith one's cognitive agency. Epistemic dependence has both a positive and a negative aspect. *Positive epistemic dependence* is when a subject exhibits a low level of cognitive agency in producing a cognitive success—i.e., of a level that would ordinarily not suffice for knowledge—but where she has knowledge nonetheless because of the influence of factors outwith her cognitive agency. *Negative epistemic dependence* is when a subject exhibits a high level of cognitive agency—i.e., of a level that would ordinarily suffice for knowledge—but where she lacks knowledge nonetheless because of the influence of factors outwith her cognitive agency.

Let's take these two forms of epistemic dependence in reverse order. The easiest way of explaining negative epistemic dependence is by appeal to an *epistemic twin earth case*. Imagine two agents who are microphysical duplicates, with identical causal histories. One of them is on earth, while the other is on twin earth. Both of them currently inhabit identical causal environments. Moreover, their *normal environment*—i.e., the kinds of things that they normal causally interact with (though perhaps not currently)—is also identical. Both agents form a true belief that *p*.

The point of keeping all these factors fixed is to ensure that everything that might conceivably have a bearing on their manifestation of cognitive agency is common to both subjects. The only difference between our two subjects concerns what is taking place in their modal environments—i.e., regarding what might have easily occurred. In particular, whereas

the subject on earth's true belief that p couldn't have easily been false, the subject on twin earth's truth belief that p could have easily been false. Here is the crux. What is taking place in one's modal environment has absolutely no bearing on one's manifestations of agency, cognitive or otherwise. Consider a non-cognitive successful manifestation of agency, such as playing the piano. The environment in which one exhibits this performance is obviously relevant to how we assess it: playing piano while underwater, for example, is obviously to manifest a very different skill than playing piano in normal conditions. But imagine now that one is playing piano in conditions where (unbeknownst to one) one could very easily have been underwater, but isn't. Is one now manifesting a different skill than is involved in playing piano in normal conditions? Surely not. As we might put the point, the successful manifestation of agency can be modally fragile, in that it is compatible with the possibility that one's success could have easily been a failure.

The problem, however, is that if we translate this point over to the cognitive case, then it follows that there will be manifestations of cognitive agency which are also modally fragile. Indeed, our epistemic twin earth example is a case in point. For although both subjects' cognitive success is equally attributable to their manifestation of cognitive ability—since we've kept fixed every conceivable factor that might potentially have a bearing in this regard—it is nonetheless the case that the subject on twin earth's cognitive success could have easily been a cognitive failure. This result is important, since knowledge is widely held to be incompatible with the modally close possibility that a belief, formed on the same basis, could be false, and for good reason. After all, to allow that knowledge could involve modally fragile cognitive success would be to allow for high levels of epistemic risk. Knowledge is distinctive in demanding not just the manifestation of a sufficient level of agency, but also the exclusion of high levels of epistemic risk—when you know, your cognitive success is both significantly attributable to your cognitive agency and also such that you couldn't have easily been wrong.[19]

The import of epistemic twin earth cases, and thus of negative epistemic dependency, is that knowledge is not simply a straightforward function of the manifestation of cognitive agency, in that one can manifest very high levels of cognitive agency and yet still fail to know because one's cognitive success is nonetheless subject to high levels of epistemic risk. In particular, notice that simply raising the threshold of cognitive agency that is required for knowledge will not solve this problem, since however high the threshold is set there will still be the potential for cases of modally fragile cognitive success.

In any case, raising the threshold would make knowledge even harder to come by than it already is by robust virtue theoretic lights. This point is particularly important in the context

of negative epistemic dependence, since this demonstrates that the virtue threshold for knowledge on this view is already too high. This because there are cases where other factors external to one's cognitive agency can play a substantive role in one's cognitive success without thereby depriving one of knowledge. The most straightforward cases in this regard concern ordinary testimonial knowledge. This is often acquired by for the most part trusting the word of another. That is, while one wouldn't just believe whatever one is told by an informant, and one wouldn't just ask any informant—i.e., such that there is some significant degree of manifestation of cognitive agency on display—in the right conditions one can nonetheless come to acquire testimonial knowledge in this fashion. The *caveat* 'in the right conditions' is important here, since what is meant by this are epistemically favourable conditions—i.e., conditions where one is not subject to high levels of epistemic risk in for the most part trusting the word of another. But still, the point remains that when one acquires knowledge in such conditions one's cognitive agency, while it plays *some* role in accounting for one's cognitive success, it is hardly the overarching element in this regard. Indeed, if anything, it is the cognitive agency of one's informant that plays the more important explanatory role.

The import of epistemic dependence, both positive and negative, is that knowledge is not simply the manifestation of a particular threshold of cognitive ability, as robust virtue epistemology claims. Rather it involves an interplay between agential and non-agential factors. Sometimes this means that in epistemically friendly conditions a subject might have knowledge even while failing to manifest high levels of cognitive agency. Sometimes this means that even despite manifesting a high level of cognitive agency, a subject nonetheless fails to know on account of some external factor.[20]

In any case, the bearing of epistemic dependence for our purposes lies in understanding how the empirical literature that is cited by proponents of epistemic situationism might in fact lend support to the idea that such cases sometimes involve positive epistemic dependence. This would mean that the subject, while having false beliefs, nonetheless is in the market for knowledge, and hence we might plausibly have a case of epistemically useful false belief, where the epistemic utility in question is indeed systematically arising out of that false belief in the required sense. In order to see this point in action, notice that many of the cognitive biases and heuristics that proponents of epistemic situationism appeal to might actually enhance, or at least not actively mitigate against, the reliability of a subject's belief-forming processes.

To take a well-worn example, subjects often make inferences from the familiarity of a place-name to how populous the town/city in question is (this is known as the *recognition heuristic*). The subject may not be aware that they are employing this heuristic, and will almost

certainly have no basis for thinking it reliable, much less will they be conscious of its limitations (i.e., the conditions under which the familiarity of a place name is not a good guide to how populous it is). Even so, this heuristic *is* a generally reliable way of forming beliefs about how populous somewhere is.[21]

Can the subject's use of this heuristic be a route to knowledge? If robust virtue epistemology were true, then this would be unlikely, since although the subject's cognitive success is gained via a reliable cognitive process, it doesn't seem to be primarily attributable to the subject's manifestation of cognitive agency. Instead, it rather seems to reflect more the good fortune that she is happening to employ a reliable cognitive heuristic. But once we have rejected robust virtue epistemology on independent grounds, and embraced the possibility of positive epistemic dependence, then we could regard such cases as constituting *bona fide* knowledge. The question would just become whether the subject exhibits *enough* cognitive ability to qualify as a knower. If so, then it wouldn't matter that some of the explanatory burden with regard to her cognitive success was carried by factors outwith her cognitive agency, such as that she happens to be employing a cognitive heuristic that is generally reliable.

Even so, have we yet got a case of epistemically useful false belief? There's certainly a case for thinking that it is epistemically beneficial to the subject to be using this heuristic 'blind' rather than trying to use it in an informed way. As Gerd Gigerenzer (e.g., 2007, 7-8) has pointed out, if one has some knowledge of the two cities that one is trying to evaluate for comparative size, then one might try to appeal to that to draw a conclusion, and this can often lead one astray. This is why, for example, foreigners can sometimes be better at making judgements about how populous a city is than people who live in the country in question, since the former have no choice but to (often subconsciously) employ the (generally reliable) cognitive heuristic. There thus might be epistemic advantages to employing these cognitive heuristics, at least when they are generally reliable.

The tricky thing, however, is to identify a particular false belief on the part of the subject which is operative here. For while it is clear that there are beliefs that the subject lacks (e.g., about how they are forming their belief in the target proposition), and while the lack of some of these beliefs might well be epistemically useful (i.e., in those cases where having more knowledge would mitigate against the use of the reliable cognitive heuristic), that's not yet to say that we have a concrete case of an epistemically useful false belief. If the subject believed that the familiarity of a place name is always a guide to how populous it is, then that would be a false belief, but why would the subject believe such a thing? In particular, insofar as she believes

anything in this regard, why would she believe this *as opposed to* believing the true proposition that this is generally reliable?

Moreover, even if she did have this false belief, why would it be epistemically useful? If anything, it seems that whether the agent believes this proposition or not is usually inert with regard to the epistemic effectiveness of the cognitive heuristic, as it's the employment of it that counts, and the agent will likely employ it either way (i.e., so long as nothing else interferes, such as coming into possession of further information that gets in the way of the employment of the heuristic). Indeed, there may be some cases where having this false belief is actually epistemically *counterproductive*, in that having this belief would prevent one from not employing the heuristic in cases where it is manifestly unreliable. A normal person wouldn't use this heuristic if one of the places in question were her own hometown, somewhere she knows to not be populous at all. In this case, clearly the familiarity of the name has everything to do with the fact that it is one's hometown and not with the size of its population. But if one really were convinced that familiarity of a town's name is always a guide to the size of this town, then one might well employ the heuristic regardless.

Of course, I've just considered one possible way in which false beliefs might be epistemically useful in the context of epistemic situationism (at least where the latter is understood in light of positive epistemic dependence). So that we have struggled to see how epistemically false beliefs might have a role to play here doesn't in any way determine that there is no scope for epistemic situationism to offer us cases of this kind. But I think we have done enough to see that the prospects in this regard are not as promising as we might have antecedently thought. In particular, it seems that it is not the subject having false beliefs that is epistemically useful in this case, but rather the subject's failure to have certain beliefs.

4. HINGE COMMITMENTS

I now want to turn to another possible case of epistemically useful false belief, which concerns the Wittgensteinian notion of a *hinge commitment*. In his very final notebooks, published as *On Certainty* (Wittgenstein 1969), Wittgenstein argued that it was essential to our rational practices—indeed, essential even for being a rational subject at all (i.e., one who can have rational beliefs, have rational doubts, and so on)—that one had certain fundamental commitments. These are the hinge commitments, about which one is optimally certain. Consider this famous passage:

> [... *T*]he *questions* that we raise and our *doubts* depend upon the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn.
>   That is to say, it belongs to the logic of our scientific investigations that certain things are *in deed* not doubted.
>   But it isn't that the situation is like this: We just *can't* investigate everything, and for that reason we are forced to rest content with assumption. If I want the door to turn, the hinges must stay put. (Wittgenstein 1969, §§341-3)

Our hinge commitments are essentially arational. Since they are required for rational evaluation to take place, they cannot themselves be subject to rational evaluation, whether positive or negative (i.e., one cannot have reasons in support of them or have reasons to doubt them). Moreover, they are neither acquired via rational processes (being instead unknowingly "swallowed down" along with other things that one is taught), nor are they responsive to rational considerations (in that the commitment would remain unaltered even if one became aware that one had no rational basis for the truth of the proposition in question). Rather, these commitments are "visceral", "animal."[22]

The notion of a hinge commitment—and how best to unpack this notion—is, of course, controversial, but I don't want to become embroiled in this controversy here.[23] My interest is rather in whether taking this idea seriously would present us with a case of epistemically useful false belief. On the face of it, the prospects look good. If it is indeed essential to being a rational subject that one has hinge commitments, it would thus follow that they have an epistemic utility, in that without them one could not be a rational inquirer at all, and hence could not even be in the market for rational belief. Moreover, this utility still stands even if the hinge commitments turn out to be false. For while in such a case one is surely unlikely to know very much—as one's beliefs will be radically in error—it is still true that having these commitments enables one to be a rational inquirer, and thus potentially enjoy at least certain kinds of rational standings (albeit non-factive ones, in contrast to knowledge). Have we identified a potential category of epistemically useful false belief?

One might initially be sceptical of this claim on the grounds that the utility in question is not epistemic but rather prudential. After all, if one really does have no rational basis for the hinge commitments, then isn't the utility in question more a matter of their practical import than anything specifically epistemic? Put another way, isn't the case for hinge commitments akin to Pascal's wager-style reasoning? Recall that Pascal's wager doesn't offer us a rational basis for thinking that God exists, but only that once one compares the utility of all the doxastic options available, belief in God is the dominant doxastic strategy. It looks like the case for hinge commitments is analogous: while we don't have a rational basis for their truth, once one

compares the doxastic options (i.e., being able to be a rational subject versus cognitive paralysis), having the hinge commitments is the dominant, and hence prudent, doxastic strategy. Crucially, however, Pascal's wager only gives us prudentially useful belief, and not epistemically useful belief, so if the analogy holds then it seems that that's all that we can derive for hinge commitments also.

The analogy does not hold, however. This is because whereas the positive utility at issue in Pascal's wager is of a purely prudential kind, the positive utility at issue with regard to hinge commitments includes positive utility which is determined from a specifically epistemic point of view. That is, while being a rational subject is certainly prudentially useful, it is also undoubtedly epistemically useful as well. Hinge commitments thus constitute an interesting case in which we have a kind of indirect epistemic rationality, in that it is epistemically rational to have these commitments even though one doesn't have a rational basis for their truth.

The real problem with thinking of hinge commitments as potential cases of epistemically useful false belief comes not from whether they are epistemically useful, but rather from the idea that they are beliefs in the specific sense of belief that we are interested in. We noted above a minimal condition for belief in this sense is that it is incompatible with agnosticism about the target proposition. A belief that *p* is a belief that *p is true*, after all. But notice that what follows from the reasoning just outlined regarding the epistemic utility of our hinge commitments doesn't provide one with any rational basis at all for the truth of the propositions in question. One has a rational basis for having a commitment to the hinges, but not for believing them to be true. It is for this reason that proponents of a Wittgensteinian epistemology who advocate that our hinge commitments are rational precisely do not characterise this commitment in terms of belief at all, but rather in terms of other propositional attitudes, like acceptance or trust, which are compatible with one being aware that one has no rational basis for the truth of the target proposition.[24] If that's right, then the epistemic utility of our hinge commitments, even when false, does not constitute a case of epistemically valuable false belief.

But I think there is a deeper concern here, which is that the reasoning just rehearsed regarding the epistemic utility of our hinge commitments betrays a faulty picture of what a hinge commitment is, at least as Wittgenstein understood that notion. Recall that for Wittgenstein our hinge commitments constitute visceral bedrock convictions on our part that are in principle unresponsive to rational considerations. It is hard to square this account of hinge commitments with a picture according to which one might endorse them on the basis of reasoning regarding their epistemic utility, much less square this account with one on which

they end up being propositional attitudes like acceptance or trusting. The second point is the more important one for our purposes, so let's focus on that.

One might accept that *p*, or trust that *p*, while being entirely agnostic about the truth of *p*. But if Wittgenstein is right then our hinge commitments are not like that at all. Indeed, the certainty that attaches to our hinge commitments precisely *excludes* the possibility that one is agnostic about their truth, where note that this certainty remains even once one recognises that one has no rational basis for their truth. Crucially, however, while hinge commitments according to Wittgenstein share with belief the property of excluding agnosticism with regard to the truth of the target proposition, they are nonetheless very different to beliefs. As noted above, the notion of belief that we are interested in bears some essential connections to rationality and truth. In particular, a basic condition for belief in this sense is that it is not the kind of propositional attitude that would remain even once one becomes aware that one has no rational basis for the truth of the target proposition. A commitment of this sort would not be a believing, but something else entirely, such as a wishful thinking. It follows that for Wittgenstein our hinge commitments, while all-out commitments to the target propositions, are not believings at all, and hence cannot be candidates for epistemically useful false beliefs.[25]

The upshot is that although our hinge commitments may well be both false and epistemically useful—and indeed epistemically useful over the long term in a systematic fashion that arises out of these very false commitments—once we are clear about the nature of the propositional attitudes in play then they are not plausible candidates for being beliefs, at least in the sense that we are interested in.

5. CONCLUDING REMARKS

Our concern has been the specific possibility of whether there can be epistemically useful false beliefs, where we have argued that this notion has to be construed in a particular fashion in order to be philosophically interesting. We have considered some initially promising cases—from the philosophy of science, from epistemic situationism, and from hinge epistemology—but have struggled to find a definitive instance. Given the piecemeal nature of our investigation, this doesn't settle the matter, since perhaps there are other regions of the epistemological realm that contain such cases. But it does at least give us a *prima facie* basis for being suspicious of the very idea that there could be a philosophically significant class of epistemically useful false beliefs.

REFERENCES

Alfano, M. (2012). Extending the Situationist Challenge to Responsibilist Virtue Epistemology. *Philosophical Quarterly* 62, 223-49.

Alfano, M. (2014). Extending the Situationist Challenge to Reliabilism About Inference. In *Virtue Scientia: Bridges Between Philosophy of Science and Virtue Epistemology*, (ed.) A. Fairweather, 103-22. Dordrecht, Holland: Springer.

Carter, J. A., & Pritchard, D. H. (*Forthcoming*). Epistemic Situationism, Epistemic Dependence, and the Epistemology of Education. In *Epistemic Situationism*, (eds.) M. Alfano & A. Fairweather. Oxford: Oxford University Press.

Davydov, V. V., & Kerr, S. T. (1995). The Influence of L. S. Vygotsky on Education Theory, Research, and Practice. *Educational Researcher* 24, 12-21.

Doris, J. (1998). Persons, Situations, and Virtue Ethics. *Noûs* 32, 504-40.

Doris, J. (2002). *Lack of Character: Personality and Moral Behavior*. Cambridge: Cambridge University Press.

Dunn, J. (2016). Epistemic Consequentialism. In *Internet Encyclopedia of Philosophy*, (eds.) B. Dowden & J. Fieser, http://www.iep.utm.edu/epis-con/.

Elgin, C. (1996). *Considered Judgment.* Princeton, NJ: Princeton University Press.

Elgin, C. (1999*a*). Education and the Advancement of Understanding. *Proceedings of the Twentieth World Congress of Philosophy* 3, 131-140.

Elgin, C. (1999*b*). Epistemology's Ends, Pedagogy's Prospects. *Facta Philosophica* 1, 39-54.

Elgin, C. (2004). True Enough. *Philosophical Issues* 14, 113-31.

Elgin, C. (2009). Is Understanding Factive? In *Epistemic Value*, (eds.) D. H. Pritchard, A. Millar & A. Haddock, 322-30. Oxford: Oxford University Press.

English, A. R. (2013). *Discontinuity in Learning. Dewey, Herbart and Education as Transformation.* Cambridge: Cambridge University Press.

Foley, J. (1994). Scaffolding. *ELT Journal* 48, 101-02.

Gigerenzer, G. (2007). *Gut Feelings: Short Cuts to Better Decision Making.* London: Penguin.

Gigerenzer, G. (2008). *Rationality for Mortals: How People Cope with Uncertainty.* Oxford: Oxford University Press.

Goldman, A. (1986) *Epistemology and Cognition.* Cambridge, MA: Harvard University Press.

Greco, J. (2009). *Achieving Knowledge.* Cambridge: Cambridge University Press.

Harman, G. (1999). Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error. *Proceedings of the Aristotelian Society* 119, 316-31.

Harman, G. (2000). The Nonexistence of Character Traits. *Proceedings of the Aristotelian Society* 100, 223-26.

Harman, G. (2003). No Character or Personality. *Business Ethics Quarterly* 13, 87-94.

Kallestrup, J., & Pritchard, D. H. (2012). Robust Virtue Epistemology and Epistemic Anti-Individualism. *Pacific Philosophical Quarterly* 93, 84-103.

Kallestrup, J., & Pritchard, D. H. (2013). Robust Virtue Epistemology and Epistemic Dependence. In *Knowledge, Virtue and Action*, (eds.) T. Henning & D. Schweikard, ch. 11. London: Routledge.

Kallestrup, J., & Pritchard, D. H. (2014). Virtue Epistemology and Epistemic Twin Earth. *European Journal of Philosophy* 22, 335-57.

MacAllister, J. (2012). Virtue Epistemology and the Philosophy of Education. *Journal of Philosophy of Education* 46, 251-70.

Plakias, A., & Tiberius, V. (2010). Well-Being. In *The Moral Psychology Handbook*, (ed.) J. Doris, 401-31. Oxford: Oxford University Press.

Pritchard, D. H. (2009*a*). Apt Performance and Epistemic Value. *Philosophical Studies* 143, 407-16.

Pritchard, D. H. (2009*b*). Knowledge, Understanding and Epistemic Value. In *Epistemology*

(*Royal Institute of Philosophy Lectures*), (ed.) A. O'Hear, 19-43. Cambridge: Cambridge University Press.

Pritchard, D. H. (2012*a*). Anti-Luck Virtue Epistemology. *Journal of Philosophy* 109, 247-79.

Pritchard, D. H. (2012*b*). Wittgenstein and the Groundlessness of Our Believing. *Synthese* 189, 255-72.

Pritchard, D. H. (2013). Epistemic Virtue and the Epistemology of Education. *Journal of Philosophy of Education* 47, 236-47.

Pritchard, D. H. (2014*a*). Re-evaluating the Epistemic Situationist Challenge to Virtue Epistemology. In *Naturalizing Epistemic Virtue*, (eds.) A. Fairweather & O. Flanagan, 143-54. Cambridge: Cambridge University Press.

Pritchard, D. H. (2014*b*). Truth as the Fundamental Epistemic Good. In *The Ethics of Belief: Individual and Social*, (eds.) J. Matheson & R. Vitz, 112-29. Oxford: Oxford University Press.

Pritchard, D. H. (2015). Wittgenstein on Hinges and Radical Scepticism in *On Certainty*. In *Blackwell Companion to Wittgenstein*, (eds.) H.-J. Glock & J. Hyman. Oxford: Blackwell.

Pritchard, D. H. (2016*a*). Epistemic *Angst. Philosophy and Phenomenological Research*, (Online First, DOI: 10.1111/phpr.12280).

Pritchard, D. H. (2016*b*). *Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing.* Princeton, NJ: Princeton University Press.

Pritchard, D. H. (2016*c*). Epistemic Dependence. *Philosophical Perspectives* 30, 1-20.

Pritchard, D. H. (2016*d*). Intellectual Virtue, Extended Cognition, and the Epistemology of Education. In *Intellectual Virtues and Education: Essays in Applied Virtue Epistemology*, (ed.) J. Baehr, 113-27. London: Routledge.

Pritchard, D. H. (*Forthcominga*). Epistemic Axiology. In *Epistemic Reasons, Epistemic Norms, and Epistemic Goals*, (eds.) M. Grajner & P. Schmechtig. Berlin: DeGruyter.

Pritchard, D. H. (*Forthcomingb*). 'Epistemic Risk', *Journal of Philosophy*.

Pritchard, D. H., Millar, A., & Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations.* Oxford: Oxford University Press.

Siegel, H. (1988). *Educating Reason: Rationality, Critical thinking, and Education*. London: Routledge.

Siegel, H. (2003). Cultivating Reason. In *Companion to the Philosophy of Education*, (ed.) R. Curren, 305-19. Oxford: Blackwell.

Simons, K. D., & Klein, J. D. (2007). The Impact of Scaffolding and Student Achievement Levels in a Problem-Based Learning Environment. *Instructional Science* 35, 41-72.

Sosa, E. (1991). *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge University Press.

Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford: Oxford University Press.

Sosa, E. (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge*. Oxford: Oxford University Press.

Sosa, E. (2016). *Judgment and Agency*. Oxford: Oxford University Press.

Stevenson, L. (2002). Six Levels of Mentality. *Philosophical Explorations* 5, 105-24.

Suarez, M. (Ed.) (2008). *Fictions in Science: Philosophical Essays on Modeling and Idealization.* London: Routledge.

Treanor, N. (2012). The Measure of Knowledge. *Noûs* 47, 557-601.

Tversky, A., & Kahneman, D. (1973). Availability: A Heuristic for Judging Frequency and Probability. *Cognitive Psychology* 5, 207-32.

Tversky, A., & Kahneman, D. (2002). Extensional versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment. In *Heuristics and Biases: The Psychology of Intuitive Judgment*, (eds.) T. Gilovich, D. Griffin & D. Kahneman, 19-48. Cambridge: Cambridge University Press.

Vygotsky, L. S. (1978). *Mind in Society: The Development of Higher Psychological Processes*. Cambridge,

MA: Harvard University Press.

Winch, C. (2006). *Education, Autonomy and Critical Thinking*. London: Routledge.

Wittgenstein, L. (1969). *On Certainty*. Edited by G. E. M. Anscombe & G. H. von Wright, translated by D. Paul & G. E. M. Anscombe. Oxford: Blackwell.

Wood, D., & Middleton, D. (1975). A Study of Assisted Problem-Solving. *British Journal of Psychology* 66, 181-91.

Wright, C. J. G. (2004). Warrant for Nothing (and Foundations for Free)? *Proceedings of the Aristotelian Society* (supp. vol.) 78, 167-212.

van Fraassen, B. (1980). *The Scientific Image,* Oxford: Clarendon Press.

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

Zagzebski, L. (1999). What is Knowledge? In *The Blackwell Guide to Epistemology*, (eds.) J. Greco & E. Sosa, 92-116. Oxford: Blackwell.

NOTES

[1] It may be a set of propositions that is believed, rather than a single proposition, but for ease of expression we will express the point in terms of a single believed proposition.

[2] See Stevenson (2002) for a helpful discussion of some of the different ways that the notion of belief is employed.

[3] Note that I will be exploring these points about belief further below. I also discuss this notion of belief in more detail in Pritchard (2016*b*, part two).

[4] Whether the empirical evidence in question really does establish this is of course controversial. For one thing, we need to be wary of how we are understanding happiness (or 'well-being', which tends to be the terminology opted for in the empirical literature). There is now a wealth of literature on well-being, and what promotes/undermines well-being. For a useful survey written from a philosophical perspective, see Plakias & Tiberius (2010).

[5] For a survey of the issues in play here, and a general defence of the idea that truth is the fundamental epistemic good, see Pritchard (2014*b*; *forthcominga*).

[6] For those not familiar with the reference, Mr. Magoo was a cartoon character who was severely shortsighted but who stubbornly refused to recognise this fact. Accordingly, he was always on the verge of one mishap or other, though as it happens he never came to any harm, albeit usually through sheer good fortune.

[7] I discuss this point in more detail in Pritchard (2014*b*; *forthcominga*).

[8] For an insightful discussion of some the issues in play here, see Treanor (2012).

[9] This is one reason why I would be sceptical about the merits of a simple-minded form of epistemic consequentialism that focused solely on the epistemic utility of belief-forming processes by assessing the number of true beliefs that they produce. See, for example, the variety of epistemic rule consequentialism defended in Goldman (1986). For a helpful overview of the recent literature on epistemic consequentialism, see Dunn (2016).

[10] For more on the idea of useful fictions in science, see the essays collected in Suarez (2008).

[11] The idea that acceptance is the appropriate propositional attitude for a scientist to take towards theoretical claims is most commonly associated with the work of van Fraassen (e.g., 1980).

[12] See, for example, Foley (1994) and Simons & Klein (2007). This idea is often traced to Vygotsky's (e.g., 1978) educational theory. See in particular his notion of the *zone of proximal development*, which effectively involves educators creating favourable learning conditions for their pupils, a process which in the contemporary educational literature is often called 'scaffolding'—e.g., Wood & Middleton (1975)—though Vygotsky never used this term himself. For a useful recent overview of Vygotsky's educational theory, see Davydov & Kerr (1995).

[13] Note that while one might naturally think of an educationally scaffolded learning environment as one that is epistemically amenable to the pupil, this needn't be so. After all, sometimes it is the role of the educator to provoke the student, and thereby disrupt her normal patterns of thought. In such cases, the student's epistemic environment will not feel very amenable from the student's point of view (though it will still be scaffolded in the relevant sense). For further discussion of this point, see English (2013).

[14] A very wise school teacher once told me that when dealing with the weakest students it was important not to correct *every* mistake they made, since the result would be an essay covered in red ink, which would obviously be dispiriting and hence ultimately counterproductive.

[15] This point relates to the wider issue of how the epistemic goals of education, properly conceived, are less geared towards the retention of information and more towards the development of higher epistemic standings closely related to the manifestation of intellectual character, such as understanding. For further discussion of the epistemic goals of education, see Siegel (1988; 2003), Elgin (1999*a*; 1999*b*), Winch (2006), MacAllister (2012), and Pritchard (2013; 2016*d*).

[16] For some of the key discussions of situationism and its relevance to virtue ethics, see Doris (1998; 2002) and Harman (1999; 2000; 2003).

[17] See especially Alfano (2012; 2014).

[18] I draw this distinction between robust and modest virtue epistemology—or 'strong' and 'weak' virtue epistemology, as I expressed the distinction in earlier work—in a number of places. See, for example, Pritchard (2009*b*; 2012*a*), Pritchard, Millar & Haddock (2010, ch. 2), and Kallestrup & Pritchard (2012; 2013; 2014).

[19] For further discussion of the notion of epistemic risk, see Pritchard (*forthcomingb*).

[20] For more on epistemic twin earth cases, and epistemic dependence more generally, see Kallestrup & Pritchard (2011; 2012; 2013) and Pritchard (2016*c*). Such cases, and the general phenomenon of epistemic dependence, build on my earlier distinction between intervening and environmental epistemic luck and the associated critique of robust virtue epistemology. See Pritchard (2009*a*; 2009*b*; 2012*a*) and Pritchard, Millar & Haddock (2010, chs. 2-4).

[21] See Gigerenzer (e.g., 2007; 2008) for a high-profile defence of the epistemic utility of some of the heuristics that we commonly employ. It is worth noting, however, that other researchers in the field often emphasise the *un*reliability of these heuristics—see, for example, Tversky & Kahneman (1973; 2002).

[22] Notice that I talk of hinge *commitments*, rather than (as is more usual in the literature), hinge *propositions*. The reason for this is that what is distinctive about Wittgenstein's conception of hinges is the particular kind of

propositional attitude in play, and not the target propositions that are subject to this attitude. As Wittgenstein notes, the very same proposition can at one time be a hinge and at another time an ordinary empirical belief.

[23] I've developed a Wittgensteinian epistemology at length elsewhere. See, especially, Pritchard (2016*b*), but also Pritchard (2012*b*; 2016*a*). For a recent survey of work on Wittgensteinian epistemology, see Pritchard (2015).

[24] See especially Wright (1994).

[25] I discuss this point at length in Pritchard (2016*b*, part 2).