Are Intellectual Virtues Truth-Relevant?

BLAKE ROEBER
University of Notre Dame

According to a view that I will call 'attributor virtue epistemology' ('AVE' for short), *S* knows that *p* only if her true belief that *p* is attributable to some intellectual virtue, competence, or ability that she possesses. AVE captures a wide range of our intuitions about the nature and value of knowledge, and it has many able defenders.[1] Unfortunately, it has an unrecognized consequence that many epistemologists will think is sufficient for rejecting it—namely, it makes knowledge depend on factors that are not *truth-relevant*, even in the broadest sense of this term, and it also makes knowledge depend in counterintuitive ways on factors that *are* truth-relevant in the more common narrow sense of this term. As I will show below, this means that the primary objection to interest-relative views in the pragmatic encroachment debate can be raised even more effectively against AVE. In §1, I draw out a consequence of AVE for a pair of cases involving perception. In §2, I use these cases to show how AVE makes knowledge depend in counterintuitive ways on truth-relevant and truth-irrelevant factors. In §3, I respond to objection. And finally, in §4, I comment on the relationship between AVE, truth-irrelevant factors, and epistemically irrelevant factors. (This article also contains an appendix, where I argue that exactly the same considerations apply to the proper functionalism views defended by Plantinga (1993, 2000), Bergmann (2006), and others.)

1. Virtuous Belief and External Agency

For the sake of simplicity, I will focus in the first instance on Sosa's version of AVE. While there are important differences between Sosa's view and other versions of AVE, it will be clear below how the lessons I draw for Sosa's view apply to any non-skeptical version of AVE.

According to Sosa (2007, 2009, 2011), knowledge is apt belief. A belief is apt just in case it's accurate because adroit. Accurate belief is just true belief, and beliefs are adroit just in case they are products of intellectual virtues, competences or skills possessed by the people who have them. Many questions arise about the *because* relation at the heart of

---

[1] Prominent attributor virtue epistemologists include Ernest Sosa (2007, 2009, 2011), John Greco (2003, 2010), Duncan Pritchard (2012), Wayne Riggs (2009), and others.

Sosa's view, but I will focus on the virtues themselves, and on the corresponding property of being adroit. Compare the following cases.

> **Actual World**: I'm a pretty normal guy, and like most people I can tell when I'm out of coffee. I just have to look and see. A moment ago, I looked to see whether my mug was empty and I saw that it was.

> **Demon World**: I'm a pretty normal guy, and like most people I can tell when I'm out of coffee. I just have to look and see. But unfortunately, I'm occasionally the victim of an evil demon. This demon is pretty weird. She loves true belief, but hates adroit belief. When she's at work, she ensures that her victims have true beliefs that aren't products of any intellectual virtues, competences, or skills that they possess. If her victims *lack* the relevant virtues, then she gives them true beliefs that they wouldn't have had if she had left them alone. But if her victims *have* the relevant virtues, she preempts those virtues. In these cases, she gives her victims exactly the true beliefs they would have had if left to their own devices, but she ensures that these beliefs are products of her own doing, rather than products of her victims' intellectual virtues. Right now, I believe that my mug is empty, just as in the actual world, and again my belief is true. But unbeknownst to me, my belief this time isn't the product of my ability to tell whether my mug is empty. A moment ago, I looked in my mug to see whether it was empty and I acquired the belief that it was. But the instant before I *would* have formed this belief on my own, the demon zapped me (or did whatever it is that evil demons do) and directly caused the belief. So while my belief is true, it's not the product of any intellectual virtue, competence or skill that I possess. Instead it's entirely the product of the demon's machinations.

The demon world is a pretty strange place, but like many strange places discussed by philosophers (Twin Earth, Swampman's bog, Mary's colorless room, *etc*.), it illustrates a point. Suppose I know that my mug is empty here in the actual world, and suppose that the demon world is identical to the actual world except for the stipulated interference of the demon. In both worlds, I believe that my mug is empty. But in the demon world, my belief can't be an instance of knowledge, if Sosa's view is true, since it's due entirely to the work of the demon, and thus not the product of any virtue, competence, or ability that I possess. This means that, on Sosa's view, I know that my mug is empty here in the actual world, but I don't know this in the demon world.

This conclusion is not an unanticipated consequence of Sosa's view. Indeed, the demon world takes its inspiration from Sosa's solution to the value problem. According to Sosa (2003), the praxical value of knowledge

> does not explain the fact that we would prefer a life of knowing, where we gain truth through our own intellectual performance, to a life where we are visited with just as much truth but through mere external agency (brainwashing, hypnosis, subliminal suggestion, *etc.*). … We do not want just truth that is given to us by happenstance, or by some alien agency, where we are given a belief that hits the mark of truth not through our own performance, not through any accomplishment creditable to us. … We want rather to attain truth by our own performance. (pp. 173-4)

On Sosa's view, knowledge is more valuable than mere true belief because knowledge is apt belief and apt belief requires that one has arrived at the truth through one's own performance, not through the work of some external agency. In the demon world, however, I arrive at the truth not through my own performance but through the work of the demon, so my belief that my mug is empty seems exactly the sort of belief that Sosa means to classify as inapt. Sosa, it seems, should be happy with the conclusion that I don't know that my mug is empty in the demon world.

Of course, Sosa's version of AVE is not the only version of AVE that entails this conclusion. According to every version of AVE, *S*'s true belief that *p* counts as knowledge only if it is attributable to some intellectual virtue, competence, or ability that she possesses. But in the demon world, I have not arrived at the truth with respect to the question whether my mug is empty because I am intellectually virtuous, competent, or skilled. The demon would have given me the same belief even if I were completely unable to tell whether my mug was empty, and my ability to tell whether my mug is empty played no role (causal or otherwise) in the actual formation of my belief.

It is important to notice that, in the demon world, I am not responding to the testimony of the demon, or anything like that. The idea isn't that the demon causes me to form the belief by whispering in my ear that my mug is empty and then letting me form the belief on the basis of her testimony, or that she causes me to form the belief by providing some stimulus that I respond to by forming the belief, in such a way that my response can be considered the product of some virtue that I possess. Instead, let's suppose, the causal path goes directly from the demon to the belief. If physicalism is true, she puts me in whatever physical state I would have been in had she left me alone and I had formed the belief on my own. If dualism is true, she puts me in whatever non-physical state I would have been in had she left me alone. This is the idea. The crucial point is that, in the demon world, my belief that my mug is empty—whatever exactly this belief amounts to (some C-fiber firing, or whatever it might be)—is not attributable to any virtue that I possess. Since AVE says that I know my mug is empty only if my belief *is* attributable to some virtue that I possess, every version of AVE entails that I don't know that my mug is empty in the demon world. And since every non-skeptical version of AVE is consistent with the stipulation that I *do* know that my mug is empty here in

the actual world, every non-skeptical version of AVE entails that I know that my mug is empty in the actual world but not in the demon world.

From here forward, by 'AVE,' I will mean *non-skeptical* versions of virtue epistemology according to which *S* knows that *p* only if her true belief that *p* is attributable to some intellectual virtue, competence, or ability that she possesses. On this usage, *every* version of AVE says that I know that my mug is empty in the actual world but not in the demon world. Is this is a problem for AVE, then? I think it is. But I don't think that the demon world gives us a counterexample to AVE. Or, at least, my argument against AVE will not rest on any intuition that, in the demon world, I *do* know that my mug is empty. Instead, I will argue that the demon world highlights certain theoretical difficulties for AVE— specifically, that it makes knowledge depend on truth-irrelevant factors in a very broad sense, and that it also makes knowledge depend in the wrong way on factors that *are* truth-relevant even in the narrow sense. The force of this objection may be opaque in this context, but it becomes clear against the backdrop of the pragmatic encroachment debate.

2. Truth-Relevance, Broad and Narrow

The pragmatic encroachment debate focuses on the question whether knowledge depends in any interesting way on our practical interests. Pragmatists answer 'yes' while purists answer 'no.'[2] The standard reason for rejecting pragmatism is that it makes knowledge depend on truth-irrelevant factors. A factor is truth-*relevant* with respect to *S*'s belief that *p* just in case it affects the probability that her belief is true, if not from her own point of view, then from some more objective perspective.[3] A factor is truth-*ir*relevant with respect to *S*'s belief that *p*, then, just in case it does not affect the probability that her belief is true, either from her own point of view, or from any more objective perspective. According to purists, it is both strongly counterintuitive and contrary to traditional epistemology that knowledge could depend on truth-irrelevant

factors. Baron Reed, for example, opens a recent argument with the following comments.

> One of the central questions that epistemologists have attempted to answer is this: when someone knows something, what explains or grounds that knowledge? Many different answers have been given—for example, clear and distinct perception, evidence, sense experience, reliable belief-producing processes—and the differences between them should not be ignored. But it is interesting to note that all of these answers agree in at least one respect: they are all focused on something that is truth directed. For example, evidence is thought to be an indication of what the truth is, reliable belief-forming processes are reliable insofar as they tend to produce true beliefs, and so forth. Traditionally, this point of agreement was so widely and deeply shared that epistemologists never really thought about it. It has been given a name—intellectualism or purism—only recently and only by the relatively small number of philosophers [*i.e.* pragmatists] who have argued against it. (2014: 95)

Similarly, Keith DeRose frames his arguments against pragmatism with these comments.

> Which features of a subject's situation are relevant to whether that subject knows some fact to be the case? … One who takes [non-evidential factors such as the reliability of a belief forming process] to be relevant to knowledge can still uphold the highly intuitive thesis that Jason Stanley has helpfully isolated and labelled 'intellectualism,' according to which the factors in virtue of which a true belief amounts to knowledge are exclusively truth-relevant, in the sense that they affect how likely it is that the belief is true, either from the point of view of the subject or from a more objective vantage point. … Recently, however, a number of epistemologists, including Stanley himself, have held that factors about a subject's situation that are not even in this broad sense truth-relevant can be relevant to whether that subject knows. (2009: 23-5)

And, as Stephen Grimm points out, even prominent *pragmatists* agree that it is counterintuitive and contrary to traditional epistemology that knowledge depends on truth-irrelevant factors. (By 'practicalism,' Grimm means pragmatism.)

> According to "orthodox" epistemology, it has recently been said, whether or not a true belief amounts to knowledge depends exclusively on truth-related factors: for example, on whether the true belief was formed in a reliable way, or was supported by good evidence, and so on. Jason Stanley (2005) refers to this as the "intellectualist" component of orthodox epistemology, and Jeremy Fantl and Matthew McGrath (2007 and forthcoming; cf. 2002) describe it as orthodox epistemology's commitment to a "purely epistemic" account of knowledge—that is, an account of knowledge where only truth-related factors figure in whether or not a person knows. … [P]racticalism is the view that whether a given true belief amounts to knowledge depends on the satisfaction of certain non-truth-related factors—in particular (it seems), it depends on whether or not the belief is appropriately responsive to the practical cost of being wrong. If its advocates are to be believed, practicalism is therefore meant to be a radical departure from traditional epistemology. If the language of *orthodoxy* is indeed apt, it is supposed to be an epistemological *heresy* of some kind. (2011, pp. 705-6, emphases Grimm's)

As these passages suggest, there is widespread agreement that, all else equal, theories of knowledge *lose* plausibility for making knowledge depend on truth-irrelevant factors.

It is worth pointing out, then, that AVE makes knowledge depend on truth-irrelevant factors. As we saw above, AVE entails that I know that my mug is empty here in the actual world but do not know this in the demon world. So notice now that there needn't be any truth-relevant difference between the actual world and the demon world with respect to my belief that my mug is empty. These differences aren't truth-relevant unless they affect the probability that my belief is true, if not from my own perspective, then at least from some other more objective epistemically relevant perspective. But we can stipulate that there is no difference between the demon world and the actual world in terms of my credence that my mug is empty, my evidence for and against the proposition that my mug is empty, the reliability of the process that produced my belief that my mug is empty, the safety of my belief that my mug is empty, and so on.[4] We don't beg any questions against AVE by building these details into the cases. On the intended understanding of both cases, the probability that my belief is true in the actual world is identical to the probability that my belief is true in the demon world, from both my perspective and from any more objective epistemically relevant perspective. But this means that there isn't any truth-relevant difference between the actual world and the demon world with respect to my belief that my mug is empty. So, if AVE is true, we have a difference in knowledge without any difference in truth-relevant factors. And this is exactly how pragmatism makes knowledge depend on truth-irrelevant factors—by generating pairs of cases that differ in terms of knowledge but not in terms of any truth-relevant factor. If purists are correct that we should reject pragmatism because it makes knowledge depend on truth-irrelevant factors, then it looks like we should reject AVE for the same reason.

This problem will strike many epistemologists as bad enough. In fact, however, AVE fairs even *worse* than pragmatism with respect to worries about the dependence of

---

[4] Of course the process responsible for my belief that my mug is empty in the demon world isn't the *same* process as the one responsible for my actual belief that my mug is empty, but this does not prevent us from stipulating that these processes are equally reliable. Notice also that, even if K = E and the proposition that mug is empty is therefore part of my total evidence here in the actual world, the proposition that my mug is empty needn't be part of my evidence *for and against the proposition that my mug is empty*. (Intuitively, the proposition that my mug is empty isn't part of my evidence for and against this very proposition even if it *is* part of my total evidence.) So, even if I know that my mug is empty in the actual world but not the demon world, the claim that K = E needn't conflict with the stipulation that my evidence for and against the proposition that my mug is empty in the demon world is identical to my actual evidence for and against the proposition that my mug is empty.

knowledge truth-relevant factors.[5] As Grimm (2011) observes, when weighing the objection that pragmatism makes knowledge depend on truth-irrelevant factors, it is important to note *how* it makes knowledge depend on truth-irrelevant factors. Consider the following familiar cases.

> **Low Stakes**: Hannah is driving past the bank on Friday afternoon. She has her paycheck in hand and she plans to deposit it, but there are long lines so she decides to deposit her paycheck on Saturday morning while she's out running errands. She knows that it does not matter much when she deposits the paycheck.

> **High Stakes**: Hannah is driving past the bank on Friday afternoon. She has her paycheck in hand and she plans to deposit it, but there are long lines so she considers returning to deposit her paycheck on Saturday morning. She believes that the bank will be open until noon on Saturday, just as in Low Stakes, but in this case she knows that there will be disastrous consequences if she does not deposit her check before noon on Saturday.[6]

According to paradigm versions of pragmatism, Hannah's belief that the bank will be open counts as knowledge in Low Stakes, but the stipulated differences in her practical interests prevent this belief from counting as knowledge in High Stakes. These stipulated differences do not take her out of position to know that the bank will be open by affecting her *evidence* for and against the proposition that the bank will be open. She has exactly the same evidence for and against this proposition in both Low Stakes and High Stakes. Instead, the stipulated differences in Hannah's practical interests take her out of position to know that the bank will be open by *raising the amount of evidence required for her to know that the bank will be open*.[7] In fact, according to these versions of pragmatism, the stipulated differences in Hannah's practical interests do not affect *any* relevant probability that Hannah's belief that the bank will be open is true. All of the relevant probabilities are the same in both Low Stakes and High Stakes. For at least the probability that the bank will be open conditional on Hannah's evidence for and against this proposition, however, they *do* affect whether it is *probable enough* that Hannah's belief that the bank will

---

[5] To put my cards on the table, I actually think it's no strike against a view that it makes knowledge depend on truth-irrelevant factors, since I think any plausible theory of knowledge will entail that knowledge depends on truth-irrelevant factors. (See Roeber 2016, for example.)

[6] See DeRose 1992, p. 913, and Stanley 2005, pp. 2-3.

[7] The standard explanation for this change usually runs in terms of the amount of evidence required for Hannah to rationally *act as if* the bank will be open, which differs from Low Stakes to High Stakes. See, for example, Ross and Schroeder 2014.

be open is true.[8] This distinction between factors that determine the probability that some belief is true and factors that determine whether it is probable enough that some belief is true is important because the latter factors (whatever they turn out to be) are clearly *epistemically* relevant even if they are not strictly speaking truth-relevant.[9]

Henceforth, let's say that a factor is truth-relevant *in the narrow sense* just in case it affects the probability that the relevant belief is true, and let's say that a factor is truth-relevant *in the broad sense* just in case it either affects the probability that the relevant belief is true, affects whether it is probable *enough* that this belief is true, or affects whether the agent has this belief in the first place. Truth-*irrelevant* factors in the broad sense will be factors that neither affect the probability that the relevant belief is true, nor affect whether it is probable enough that this belief is true, nor affect whether the agent has the belief. As Grimm points out (ibid), paradigm versions of pragmatism make knowledge depend on truth-irrelevant factors in the narrow sense, but they don't make knowledge depend on truth-irrelevant factors in the broad sense, since they entail that practical interests will only affect knowledge when they affect whether it is probable enough that the belief in question is true, not by affecting the probability that this belief is true, but instead by moving some relevant threshold.

While pragmatism makes knowledge depend on factors that aren't truth-relevant in the narrow sense, AVE makes knowledge depend on factors that aren't truth-relevant even in the broad sense. Again, compare my actual belief that my mug is empty to my belief that my mug is empty in the demon world. As we have already seen, there needn't be any difference between the demon world and the actual world in terms of my credence that my mug is empty, my evidence that my mug is empty, the reliability of the process that produced my belief that my mug is empty, the safety of my belief that my mug is empty, and so on. So now notice that there needn't be any difference between the demon world and the actual world in terms of any relevant thresholds, either. The minimum credence consistent with my knowing that my mug is empty, the minimum amount of evidence required for knowing that my mug is empty, the minimum levels of reliability or safety consistent with my knowing this, *etc.*, might be exactly as high in the demon world as they are in the actual world. Thus, for any relevant probability, the minimum probability that is consistent with knowledge might be exactly as high in the demon world as it is in the actual world. Indeed, since the demon world and the actual world only differ with respect to the stipulated activity of the demon, the demon world

---

[8] By "it is probable enough that Hannah's belief is true," I mean that the probability that her belief is true (whatever exactly it is) does not entail that Hannah does not know that the bank will be open. Her belief satisfies whatever probabilistic requirements there are on knowledge.

[9] See, for example, Brown 2014.

and the actual world are presumably *identical* with respect to whatever factors set the minimum relevant probabilities consistent with knowledge. But as we saw above, if any version of AVE is true, then I know that my mug is empty here in the actual world but not in the demon world. The upshot is that, while pragmatism only makes knowledge depend on factors that aren't truth-relevant in the narrow sense, AVE makes knowledge depend on factors that aren't truth-relevant even in the broad sense, since it makes knowledge depend on factors that affect neither the probability that the belief in question is true, nor whether it is probable enough that this belief is true, nor whether I hold this belief in the first place.

In fact, we can press these worries about truth-relevance even further. There are only two ways a factor can affect the probability that my belief is true: positively or negatively—positively just in case it *raises* the probability that my belief is true, and negatively just in case it *lowers* the probability that my belief is true. Truth-irrelevant factors have *no* affect on the probability that a belief is true. If it's hard to believe that truth-irrelevant factors can disqualify my belief from counting as knowledge, it's surely much harder to believe that factors that only *positively* affect the probability that my belief is true can disqualify it from counting as knowledge. But this is exactly what AVE entails. Just imagine that the demon never gives anyone a false belief, so that my belief in the demon world is the product of a perfectly reliable, completely infallible process. Then, if AVE is true, there are facts about my belief in the demon world that make it literally *certain* that my belief is true, and yet these facts disqualify my belief from counting as knowledge. But intuitively, the relationship between knowledge and truth goes the other way around. Intuitively, factors that make it certain that my belief is true should only count in *favor* of my belief's being an instance of knowledge. So, it seems, AVE not only makes knowledge depend on factors that aren't truth-relevant in even the broad sense, it also makes knowledge depend *in the wrong way* on factors that *are* truth-relevant in the narrow sense.

3. Belief, Agency, and the Scope of AVE

At this point, readers sympathetic to Sosa's idea that belief is a species of performance (2011: 1), or sympathetic to his distinction between functional and judgmental belief and his subsequent emphasis on the latter (2015: 51-2), might be tempted to respond to the arguments in §2 in one of two ways. First, they might be tempted to say that, because a mental state which isn't the product of my own agency can't be one of *my* beliefs (or one of my *beliefs*), I lack knowledge in the demon world for the mundane reason that I lack the relevant belief in the demon world. Second, they might be tempted to restrict AVE

to judgemental beliefs (beliefs that *are* products of my own agency), so that AVE doesn't say anything about my belief in the demon world.

Both responses avoid the conclusions of §2, but neither response is very satisfying. Of course, Sosa himself would not agree with the first response, since he thinks that "brainwashing, hypnosis, subliminal suggestion, *etc.*," might give us beliefs that aren't products of our own agency (2003: 173). More importantly, the first response saddles its proponent with an implausible theory of mind. By hypothesis, if physicalism is true, the evil demon puts me in whatever physical state I would have been in had she left me alone and I had formed the belief that my mug is empty on my own, in exactly the way that I normally form such beliefs. And by hypothesis, if dualism is true, then she puts me in whatever non-physical state I would have been in had she left me alone and I had formed this belief the regular way. By hypothesis, there is no difference between the demon world and the actual world in terms of any of my dispositions. By hypothesis, I am disposed to assert that my mug is empty here in the actual world iff I am disposed to assert this in the demon world, I am disposed to affirm that my mug is empty here in the actual world iff I am disposed to affirm this in the demon world, I am disposed to use the proposition that my mug is empty as a premise in practical or theoretical reasoning here in the actual world iff I am disposed to do this in the demon world, and so on. Nor is there is any difference between the demon world and the actual world in terms of any of my desires, any of my intentions, any of my actions, any of my inferences, or anything like that. By hypothesis, if I desire to get more coffee here in the actual world, form the intention to get more coffee, and then walk over to the pot, I do all of this in the demon world too. Whatever test either of us might perform to see if I believe that my mug is empty here in the actual world, I would pass this test in exactly the same way in the demon world. A truth serum wouldn't reveal any difference between the demon world and the actual world. Nor would picking up my mug and turning it upside down over my computer, or anything like that. A CAT scan would show exactly the same thing in the demon world as it would show here in the actual world, and so on. This is all true on the intended understanding of the cases. Given all this, however, it seems implausible that I don't believe that my mug is empty in the demon world simply because, in the demon world, my mental state is not the product of my own agency. So, it seems doubtful that the first reply will work.

The obvious problem with the second reply, however, is that it just changes the subject. Granted, the demon world doesn't cause any problem for the claim that, *if S*'s belief that *p* is the product of her own agency, *then* she knows that *p* only if her belief that *p* is attributable to some intellectual virtue that she possesses. After all, this claim doesn't say anything at all about my belief that my mug is empty in the demon world. But this is

not the claim that I am arguing against. I am arguing against the claim that *S* knows that *p* only if her belief that *p* is attributable to some intellectual virtue that she possesses, and the demon world *does* cause trouble for this claim. Since, by 'AVE,' I just mean non-skeptical versions of virtue epistemology according to which *S* knows that *p* only if her belief that *p* is attributable to some intellectual virtue she possesses, the demon world causes trouble for AVE. This conclusion would not be very interesting if AVE were just a straw person and it had no defenders, of course. But as I pointed out in the opening paragraph of this paper, AVE has many prominent defenders. So neither objection works.

At this point, attributor virtue epistemologists might reply that it's just intuitively obvious that I don't know that my mug is empty in the demon world, and that the demon world just shows that knowledge is stranger than many of us thought. If it *is* intuitively obvious that I don't know that my mug is empty in the demon world, this reply is exactly right. But it's an empirical question whether people have the intuition that I lack knowledge in the demon world, and nobody has done the relevant study. And in any case, I lack this intuition, since the case doesn't strike me one way or another. So at least at this point, attributor virtue epistemologists can't respond to the arguments in §2 by saying that it is intuitively obvious that I lack knowledge in the demon world.

4. Conclusion

While many contributors to the virtue epistemology debate also contribute to the pragmatic encroachment debate (and *vice versa*), nobody has noticed that the primary objection to pragmatism can be pushed even more forcefully against AVE. Why is this? Perhaps because, in many of the central cases discussed in the virtue epistemology literature, people are concerned with the presence or absence of the relevant virtues (not with whether some virtues that *are* present are causally responsible for the relevant belief), and the presence or absence of the relevant virtues typically makes exactly the truth-relevant difference that virtue epistemologists would want it to make. Take Jennifer Lackey's "Chicago Visitor" case, for example.

> **Chicago Visitor:** Having just arrived at the train station in Chicago, Morris wishes to obtain directions to the Sears Tower. He looks around, approaches the first adult passerby that he sees, and asks how to get to his desired destination. The passerby, who happens to be a lifelong resident of Chicago and knows the city extraordinarily well, provides Morris with impeccable directions to the Sears Tower by telling him that it is located two blocks east of the train station. Morris forms the corresponding true belief. (2009: 29)

Lackey presents Chicago Visitor as part of a dilemma for a species of virtue epistemology that she calls "the credit view." According to the credit view, *S* knows that *p* only if she deserves credit for arriving at the true belief that *p*.[10] The challenge for the credit view is that, while Morris seems to know where the Sears Tower is located, it's the *passerby* who seems to deserve all the credit for Morris's arriving at the true belief that the Sears Tower is two blocks east of the train station.

In response to Chicago Visitor, credit theorists customarily argue that, unless we assume that Morris's belief results from some ability to discriminate reliable from unreliable sources of information, we will have no reason to doubt (*e.g.*) that Morris could easily have asked a lost but overconfident tourist for directions instead of the reliable local that he happened to ask. But (the response continues), if he could just as easily have asked a lost tourist, it's clear that his belief that the Sears Tower is two blocks east of the train station falls short of knowledge. Since Morris deserves at least partial credit for his true belief if his belief *has* resulted from some ability to discriminate reliable from unreliable sources of information (the response concludes), either Morris *is* credit-worthy, or we have no reason to think that his belief counts as knowledge. Either way, we lack a clear case of knowledge without credit.[11] According to Lackey, this response leads to trouble elsewhere.[12] But notice that this response appeals to *truth-relevant* differences between Chicago Visitor and the case where Morris could easily have asked a lost but overconfident tourist for directions. After all, the probability that Morris arrives at the true belief that the Sears Tower is two blocks east of the train station is much higher in cases where Morris can discriminate reliable from unreliable sources of information than in cases where it's only by dumb luck that Morris asks a reliable local rather than a lost but overconfident tourist.

Since we must consider bizarre scenarios like the demon world to see how intellectual virtues might fail to be truth-relevant, it is perhaps unsurprising that even those epistemologists who contribute to both the pragmatic encroachment debate and the virtue epistemology debate have failed to notice that AVE makes knowledge depend on truth-irrelevant factors. Places like the demon world show that the primary objection to

---

[10] The credit view presumably entails AVE since, presumably, *S* deserves credit for her true belief that *p* only if her belief is the product of some intellectual virtue that she possess. If this is right, then everything I said above about AVE also applies to the credit view.

[11] See, for example, Riggs 2009 (p. 209), Greco 2007 (p. 63), and Sosa 2007 (p. 95).

[12] Lackey uses Chicago Visitor as part of a *dilemma* for the credit view rather than simply arguing that Chicago Visitor constitutes a straightforward counterexample to the credit view. As she argues, if we make the relevant notion of credit weak enough to count Morris as knowing where the Sears Tower is located, the credit view will have trouble denying knowledge in Gettier cases like Ginet's fake barn case. Thus, concludes Lackey, the credit view either gets the wrong result in cases like Chicago Visitor, or it gets the wrong result in cases like Ginet's fake barn case.

interest-relative views in the pragmatic encroachment debate can be raised even more effectively against AVE. Since the vast majority of purported necessary conditions on knowledge are truth-relevant in at least the broad sense, it's tempting to conclude that AVE makes knowledge depend not just on *truth*-irrelevant factors, but on fully *epistemically* irrelevant factors. To put my cards on the table, however, AVE illuminates so much in epistemology that I still find it attractive—perhaps even more so than any competing view. Since affecting whether a belief counts as knowledge seems sufficient for being epistemically irrelevant, I am uncomfortable saying that AVE makes knowledge depend on epistemically irrelevant factors. But in any case, the fact that AVE makes knowledge depend on factors that aren't truth-relevant even in the broad sense, and the way that it also makes knowledge depend in counterintuitive ways on factors that *are* truth-relevant in the narrow sense, clearly lend at least some credibility to alternative views.

Appendix: Proper Functionalism

According to the proper-functionalist views defended by Plantinga (1993, 2000), Bergmann (2006), and others, *S* knows that *p* only if her belief that *p* is the product of properly functioning cognitive faculties. Proper functionalism makes knowledge depend on truth-irrelevant factors, and depend in counterintuitive ways on factors that are truth-relevant, in exactly the way that AVE makes knowledge depend on truth-irrelevant factors and depend in counterintuitive ways on factors that are truth-relevant.

We don't beg any question against proper functionalism by stipulating that I know that my mug is empty here in the actual world. In the demon world, however, my belief that my mug is empty is no more the product of properly functioning cognitive faculties than my pulse would be the product of properly functioning organs if a paramedic who mistakenly thought I was having a heart attack preempted the normal beating of my heart by performing compressions on my chest. A world where the demon loves true beliefs but hates beliefs produced by properly functioning cognitive faculties would have exactly the same implications for proper functionalism as the world where the demon loves true beliefs but hates adroit beliefs has for AVE. If proper functionalism is true, the actual world and this demon world would give us a pair of cases that differ with respect to my knowledge even though they are identical on every truth-relevant dimension, even in the broad sense of this term. And a world where the demon never gives anyone a false belief would give us a case where I lose my knowledge in virtue of a truth-relevant difference that only makes it *more* likely that my belief is true, since my actual belief is the product of fallible cognitive faculties. So demon worlds where the

demon hates beliefs produced by properly functioning cognitive faculties create exactly the same trouble for proper functionalism as demon worlds where the demon hates adroit beliefs create for AVE.[13]

References

Bergmann, Michael (2006). *Justification without Awareness* (Oxford: Oxford University Press).

Brown, Jessica (2008). "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning," *Noûs*, Vol. 42, No. 2: pp. 167-89.

——— (2012). "Practical Reasoning, Decision Theory, and Anti-Intellectualism," *Episteme*, Vol. 9, No. 1: pp. 43-62.

——— (2013). "Experimental Philosophy, Contextualism, and SSI," *Philosophy and Phenomenological Research*, Vol. 86, No. 2: pp. 233-61.

——— (2014). "Impurism, Practical Reasoning, and the Threshold Problem," *Noûs*, Vol. 48, No. 1: pp. 179-92.

Cohen, Stewart (2012). "Does Practical Rationality Constrain Epistemic Rationality?," *Philosophy and Phenomenological Research*, Vol. 85, No. 2.

DeRose, Keith (2009). *The Case for Contextualism: Volume 1* (Oxford: Oxford University Press).

Fantl, Jeremy and Matthew McGrath (2009). *Knowledge in an Uncertain World* (Oxford: Oxford University Press).

Fumerton, Richard (2010). "Fencing Out Pragmatic Encroachment," in *Philosophical Perspectives*, Volume 24, Epistemology: pp. 243-53.

Ganson, Dorit (2008). "Evidentialism and Pragmatic Constraints on Outright Belief," *Philosophical Studies*, Vol. 139: pp. 441-58.

Greco, John (2003). "Knowledge as Credit for True Belief," in *Intellectual Virtue: Perspectives from Ethics to Epistemology*, M. DePaul and L. Zagzebski Eds. (New York: Oxford University Press).

——— (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*, (Cambridge: Cambridge University Press).

Grimm, Stephen (2011). "On Intellectualism in Epistemology," *Mind*, Vol. 120, No. 479: pp. 705-33.

Lackey, Jennifer (2009). "Knowledge and Credit," *Philosophical Studies*, Vol. 142: pp. 27-42.

Nagel, Jennifer (2008). "Knowledge Ascriptions and the Psychological Consequences of Changing Stakes," *Australasian Journal of Philosophy*, Vol. 86, No. 2: 279-93.

——— (2010). "Epistemic Anxiety and Adaptive Invariantism," *Philosophical Perspectives*, Vol. 24: pp. 407-35.

Neta, Ram (2007). "Anti-intellectualism and the Knowledge-Action Principle," *Philosophy and Phenomenological Research*, Vol. 75, No. 1: pp. 180-7.

——— (2012). "The Case Against Purity," *Philosophy and Phenomenological Research*, Vol. 85, No. 2.

Plantinga, Alvin (1993). *Warrant and Proper Function* (Oxford: Oxford University Press).

——— (2000). *Warranted Christian Belief* (Oxford: Oxford University Press).

Pritchard, Duncan (2012). "Anti-Luck Virtue Epistemology," *The Journal of Philosophy*, Vol. 109, No. 3: pp. 247-79.

Reed, Baron (2014). "Practical Matters Do Not Affect Whether You Know," in *Contemporary Debates in Epistemology: Second Edition*, M. Steup, J. Turri, and E. Sosa Eds. (Oxford: Wiley Blackwell Publishing).

Riggs, Wayne (2009), "Two Problems of Easy Credit," *Synthese*, Vol. 169, No. 1: pp. 201-216.

Roeber, Blake (forthcoming), "Anti-Intellectualism," *Mind*.

Ross, Jacob and Mark Schroeder (2014). "Belief, Credence and Pragmatic Encroachment," *Philosophy and Phenomenological Research*, Vol. 88, No. 2: pp. 259-88.

Schroeder, Mark (2012). "Stakes, Withholding, and Pragmatic Encroachment on Knowledge," *Philosophical Studies*, Vol. 160, No. 2: 265-85.

Stanley, Jason (2005). *Knowledge and Practical Interests* (Oxford: Oxford University Press).

Sosa, Ernest (2003). "The Place of Truth in Epistemology," in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, M. DePaul and L. Zagzebski Eds. (Oxford: Oxford University press).

_____ (2007). *A Virtue Epistemology* (Oxford: Oxford University Press).

_____ (2009). *Reflective Knowledge* (Oxford: Oxford University Press).

_____ (2011). *Knowing Full Well* (Princeton: Princeton University Press).

_____ (2015). *Judgment and Agency* (Oxford: Oxford University Press).

Weatherson, Brian (2005). "Can We Do Without Pragmatic Encroachment," *Philosophical Perspectives*, Vol. 19: pp. 418-43.

_____ (2012). "Knowledge, Bets and Interests," in *New Essays on Knowledge Ascriptions*, eds. J. Brown and M. Gerken (Oxford: Oxford University Press).