

Competent Perspectives and the New Evil Demon Problem

Lisa Miracchi
University of Pennsylvania

December 20, 2015

Forthcoming in *The New Evil Demon: New Essays on Knowledge, Justification and Rationality*, Oxford University Press, eds. Fabian Dorsch and Julien Dutant.

The New Evil Demon problem is a problem for externalist theories of justification, and has been a subject of ongoing debate since it was introduced in 1983 by Keith Lehrer and Stewart Cohen.¹ Lehrer and Cohen ask us to:

Imagine that, unknown to us, our cognitive processes, those involved in perception, memory and inference, are rendered unreliable by the actions of a powerful demon or malevolent scientist. It would follow on reliabilist views that under such conditions the beliefs generated by those processes would not be justified. This result is unacceptable. The truth of the demon hypothesis also entails that our experiences and our reasonings are just what they would be if our cognitive processes were reliable, and, therefore, that we would be just as well justified in believing what we do if the demon hypothesis were true as if it were false.

—Lehrer & Cohen (1983), p. 192

In other words, Lehrer and Cohen ask us to imagine that the subjective, first-personal character of our mental lives is the same as it normally is, but that we are being radically deceived, so that (nearly?) none of our beliefs are connected to the world in the way they normally are, in the way that provides us with knowledge of the world.² Nevertheless, Lehrer and Cohen maintain, such a deceived subject would still be justified. Let us call duplicates of the sort Lehrer and Cohen are imagining *perspectival duplicates*.

¹Thanks to Nic Bommarito, Cameron Boult, Brian Cutter, Julien Dutant, John Greco, Christoph Kelp, Rachel McKinney, Alan Millar, Ernest Sosa, Kurt Sylvan, and Alex Worsnip.

²See also Cohen (1984).

The internalism/ externalism distinction is made in many different ways, but one way to make it is to distinguish those who think that disconnection from the world in the evil demon scenario precludes justification (externalists) from those who do not (internalists). Externalists hold that one's reliable connection to the world is all that matters for justification; internalists hold that what matters is the subject's perspective.

Many externalists reject the internalist's claim that subjects in the evil demon scenario have a positive epistemic standing in common with normal subjects, maintaining that positive epistemic standing is entirely a matter of how one is connected to the world. This, however, has well-known problems. For example, let us start with someone who is clearly epistemically virtuous—someone who seeks out proper evidence on questions of interest, reasons thoroughly and effectively about these questions, and persistently develops and hones her intellectual abilities in the service of acquiring and maintaining knowledge of herself and the world. Let us also take someone who is clearly epistemically vicious—someone who is intellectually lazy, who indulges in wishful thinking and bad reasoning that supports whatever she wants to believe, who refuses to acquire new ways of thinking or reasoning despite having good evidence that her current methods are misleading, and so on.

Now, we imagine that perspectival duplicates of our two subjects are in evil demon scenarios. Are these duplicates epistemically on a par with each other? Intuitively not: unfortunate circumstances make both of their belief-forming methods equally completely unreliable, but that does not not erase all epistemic differences.³ The duplicate of our virtuous epistemic agent is still more virtuous than the duplicate of our vicious epistemic agent, for she is still reasoning better. Internalists take this kind of consequence to be a *reductio* of the externalist position—epistemic standing must be the sort of thing that differentiates between these two perspectival duplicates.⁴

Internalists are in large part concerned with trying to capture the difference that our perspective on the world and on our own abilities makes

³Here and throughout I will ask whether the beliefs of an epistemic subject are justified/rational or not and why. This is not meant to be read as a question about *all* the beliefs of a subject. I assume that virtuous duplicates (and agents) are capable of failing to believe from epistemic competence from time to time, and these cases are not of interest here. Likewise, *mutatis mutandis* for other agents. We are considering here the epistemic status of beliefs that are exercises of epistemic competences or other propensities to believe, and it is only for the sake of concision that I talk about “the beliefs” of certain kinds of epistemic agents.

⁴See Cohen (1984), p. 283 for discussion of this kind of case.

to our epistemic standing. Without it, they worry, our epistemic standing becomes too divorced from what makes us human, too much a matter of a mere machine functioning properly, rather than a person grappling with questions. I agree. An adequate theory of our epistemic standing must take into account our subjective perspectives on the world, how things seem to us, what considerations we are bearing in mind, how we are trying to reason, and whether we are properly committed to knowing how the world is, even if it is not always in our favor.

However, the traditional internalist account of justification has its serious shortfalls too. It fails to capture the sense in which our virtuous epistemic agent is epistemically better off than her perspectival counterpart—for one, she is actually competent. She doesn't just mean well, she reliably and effectively acquires and maintains knowledge about her environment. Even when she goes wrong and falsely believes, her errors are of an entirely different magnitude than the errors her perspectival duplicate makes. This difference should be reflected in our account of epistemic standing.

One strategy for solving this problem which has seemed attractive to many is to allow for more than just one kind of positive epistemic standing that falls short of knowledge. Perhaps we should separate the externalist notion of justification (what I'll henceforth call "justification") from the internalist notion of justification ("rationality"), letting everyone agree to disagree. As things stand, however, this move has serious costs. First, many internalists and externalists hold that knowledge can be analyzed in terms of justification, truth, and some anti-Gettier feature.⁵ For such views the question arises: which kind of justification should serve in such an analysis? If it is the externalist kind, why does the internalist kind also seem necessary for knowledge? What work is it doing in our epistemology then? (And vice versa.) What do these two kinds of epistemic standing have to do with each other? Why do justification and rationality both count as positive *epistemic* statuses? Are there more positive epistemic statuses? Why stop at these two? Why not let a thousand flowers bloom?

One might think these issues are avoided by those adopting a knowledge-first approach, but at least initially the problem becomes even worse. Knowledge-firsters claim that knowledge is not analyzable in terms of justification (or rationality) but is rather epistemically fundamental, and all epistemic statuses are derivative from the epistemic status of knowledge.⁶ But now we

⁵This project clearly motivates both Lehrer and Cohen.

⁶Knowledge-firsters differ on whether they claim knowledge is conceptually or metaphysically unanalyzable in terms of other epistemic and mental features. Here and else-

have not one but *two* kinds of positive epistemic status that need accounting for in terms of knowledge. Knowledge-firsters either have to settle for a less explanatory epistemic theory than belief-firsters or they have their work cut out for them. Being a knowledge-firster myself, this much mystery makes me pretty uneasy.

Whew. This is where the debate is right now, and it is no wonder that many externalists (both knowledge-firsters and belief-firsters) have turned to the idea of an *excuse* to try to get out of the trouble.⁷ On this view, only beliefs that are knowledge have positive epistemic status at all, but there are other cases where a subject may be blameless, or excusable, for having a belief. Evil demon scenarios, they claim, are cases of this sort. However, this strategy doesn't fix the problem.

First, it doesn't explain *why* the subject is epistemically excusable in virtue of being a perspectival duplicate. What is it about our perspectives that makes a difference to blameworthiness?⁸ Often this is assumed, but it cannot be in this context. Here we return to the original internalist call to make our humanity relevant to epistemology—to make our interests and goals have a role to play in our epistemic lives. What is it, then, about our perspectives that provides us with excuses? Moreover, why do our first-personal perspectives make a difference to excusability despite not making a difference to genuine epistemic standing? More needs to be said here than has been done to date.⁹

where I am primarily interested in metaphysical questions. See Ichikawa & Jenkins (manuscript) for a helpful discussion of the diversity of knowledge-first views.

⁷Williamson (this volume), Littlejohn (2012).

⁸E.g. Littlejohn (2012) argues that evil demon subjects still “pursue their epistemic ends rationally and responsibly” (59) and that this provides them with an excuse for believing as they do. Why is it, however, that evil demon subjects have these virtues despite being wholly unreliable, and why should these provide one with an excuse for not having genuine justification? As Littlejohn himself agrees (58), the strategy treating our world as the normal world for *any* subject in *any* world (as Williamson (this volume), Comesaña (2002), and others do), is an undue privileging of our own situation rather than a genuine explanation of the norms applicable in the evil demon world.

⁹As my main aim in this paper is to offer my own solution, I cannot defend this claim in detail. However, to better see the kind of worries at issue, I will briefly consider Williamson (this volume)'s proposed solution to the new evil demon problem by appealing to derivative norms. He argues that although evil demon subjects violate the primary epistemic norm—believe only what you know—they can satisfy the secondary and tertiary norms of having a general disposition to believe only what one knows and doing what a person who had such a disposition would do in such a situation. Satisfying these norms, he claims, is sufficient for having an excuse for violating the primary norm. He does not, however, explain (i) *why* certain secondary and tertiary norms are generated by a primary norm, (ii) *why* the virtuous duplicate complies with these derivative norms, given that her world is thoroughly unlike

Second, proponents of this strategy also use it to account for cases in which we have a justified false belief in normal environments.¹⁰ They thus fail to account for the epistemic difference between our virtuous agent who exercises her epistemic competence and yet believes falsely on a particular occasion, and her perspectival duplicate. Lastly and most importantly, this kind of strategy fails to explain the way in which the perspectival duplicate of our virtuous epistemic agent is in some sense also *virtuous*. It is not merely that she is less in the wrong than the vicious duplicate; she is doing *something* epistemically right, at least, and we should be able to give an account of what it is.¹¹

It's time for a new strategy. In what follows, I will extend my direct virtue epistemology (2015a; 2015b) to explain how a knowledge-first framework can account for two kinds of positive epistemic standing, one tracked by externalists, who claim that the virtuous duplicate lacks justification, the other tracked by internalists, who claim that the virtuous duplicate has justification, and moreover that such justification is not enjoyed by the vicious duplicate. It also explains what these kinds of epistemic standing have to do with each other. In short, I will argue that all justified beliefs are *good candidates* for knowledge, and are such because they are exercises of competences to know. However, there are two importantly different senses in which a belief may be a good candidate for knowledge, one corresponding to an externalist kind of justification and the other corresponding to an internalist one.

In section 1, I discuss the New Evil Demon problem in more depth, and argue that externalists cannot easily dismiss it. In section 2, I review some core features of my direct virtue epistemology and explain how it already delivers an externalist kind of justification. In section 3, I explain what kind of positive epistemic standing perspectival duplicates have, and why this epistemic standing is dependent on the normative status of knowledge. In section 4, I show how this normative status may be explained using the

ours (he just claims that our world is the world of evaluation for normal scenarios (p. 14)), and (iii) why compliance with the derivative norms should provide one with an excuse for violating a primary norm. Such a task is necessary for a full account of epistemic excuse and seems just as difficult as that of explaining how there can be two kinds of positive epistemic status.

¹⁰See, e.g. Williamson (this volume).

¹¹Some sophisticated versions of the view might have something to say here. E.g., Littlejohn (2012) admits that there is a sense of justification that corresponds to the internalist sense of justification (personal justification) and claims that evil demon subjects can have it. However, we want a unified account of epistemic standing, and Littlejohn does not provide one.

tools of virtue epistemology. In section 5, I show how the account solves the new evil demon problem in a more satisfactory way than existing accounts. We end up with a view of knowledge, justification, and rationality that is plausible, motivated, and theoretically unified.

1 What Exactly Is the New Evil Demon Problem?

The New Evil Demon Problem asks us to consider a scenario in which everything seems to be the same to us as it normally does, but in which we are radically deceived by some evil creature with the power to make us undergo such a persistent illusion. For simplicity's sake, I will consider the way of filling out the case on which we are, and always have been, brains in vats which an evil demon has made to have the thoughts, experiences, etc., we would have if things were normal.¹²

A normal question for someone with externalist leanings to have at the very outset is whether such a scenario is even possible. Sure, I can imagine being a brain in a vat and having all the same experiences and thoughts as I do now, but could I actually *be one*?

An influential line of argument says *No*. If our experiences, thoughts, and so on are in part determined by our relations to our environment, then we couldn't have the same experiences, thoughts, and so on if we were brains in vats. Such considerations are typically brought up to counter skeptical worries about our knowledge of the world.¹³ We appeal to the idea that what our mental states are about is at least in part determined by our relations to the world in order to rule out the possibility that we could be radically deceived.¹⁴

Adopting inspiration from this approach, the hardline externalist about justification might then push the point here. If I couldn't experience or think about the same things in the evil demon scenario as I do now, then I wouldn't have the same experiences or thoughts. Things wouldn't seem *to be* the same to me. Thus the brain in a vat scenario is not one in which I form the same beliefs on the basis of the same experiences in the same ways but now these ways are highly unreliable. The new evil demon scenario is

¹²This construal ignores complications that arise from recent envatment. These complications won't change the moral of the story.

¹³See Putnam (1981).

¹⁴There are of course questions to ask about this strategy (e.g. see Brueckner (1986) for plausible worries), but the point here is not to show that the externalist has a convincing response to skepticism. Rather, I just wish to point out that it is a move someone with externalist leanings (both in semantics and in epistemology) might plausibly make.

thus metaphysically impossible, and so not a counterexample to externalism. We can happily disregard troubling intuitions about such cases.

I think that this response to the internalist's challenge misses the point. All that matters, in order for the evil demon scenario to pose a problem for externalism about justification, is for it to be *conceivable*, not possible. The point of raising the scenario is not that externalism fails to be extensionally adequate, but that it wholly credits reliability with being responsible for epistemic standing. The scenario is being used to teach us that beliefs have certain kinds of epistemic standing in virtue of certain properties of our perspectival lives—if our perspectival lives could be preserved while our reliable connection to the world were severed, some things would still be going epistemically right with us. The virtuous perspectival duplicate would still be reasoning in the right sort of way—she would believe properly on the basis of her experience, she would engage in proper inferences, ask the right questions, and so on. She would have a certain epistemic standing that is preserved because it is determined by certain perspectival features of her mental life. It is irrelevant whether one could, as a matter of metaphysical possibility, have this aspect of one's mental life preserved in the absence of reliable connections to the environment. What matters for the epistemic internalist is that any contributions that reliable connections to the environment make to epistemic standing are made *via* their giving rise to the perspectival aspects of our mental lives.

This is a point that is often overlooked and so it is worth making again: sometimes, all you need in order for a case to make a point is for it to conceptually separate two properties. This can show that intuitively certain properties *A* are responsible for epistemic standing *N* (or whatever philosophically interesting feature you're interested in) and other properties *B* are not. It doesn't matter whether the *A* properties could exist without the *B* properties. Perhaps the *B* properties are metaphysically necessary for the *A* properties. But even if that is so, if our intuitions are on the right track, the *B* properties contribute to *N* only *via* grounding the *A* properties.

Accordingly, the internalist can claim that content externalism is beside the point. Perhaps there couldn't be radical deception scenarios where I form the belief that I'm sitting under a tree on the basis of my experience of doing so. Nevertheless, when we imagine the case, we judge that the subject is doing something right, whereas someone who judges that there are pink elephants in the room on the basis of the same experience is not doing something right. *That* is all we need to suppose in order for the evil demon scenario to generate a problem for externalism.

A solution to the new evil demon problem, then, will explain how there

is a kind of epistemic standing that is not *directly* determined by our reliable hook-up to the world, but rather by our mental, first-personal, perspectives: how things seem to us, how we are reasoning, what we are aiming at when we are reasoning in certain ways or asking certain questions. However, a solution need not appeal to only features that are present in the evil demon scenario. If the scenario is, as seems highly plausible to me, metaphysically impossible, then our account of which epistemic agents have the perspectival features responsible for rationality may indeed appeal to how the subject is related to the world in normal cases. We should accept the internalist point that mental features make a direct difference to epistemic standing without conceding that such mental features are solely determined by what is inside the head.

That the mental features responsible for internalist justification *do* depend on what is in the world can be illustrated by introducing a third perspectival duplicate: a merely well-meaning one. A *merely well-meaning* agent values knowledge and tries to form and maintain beliefs in knowledgeable ways, but systematically and widely fails. She doesn't have a good sense of what considerations bear on questions of interest; her reasoning is not logical, or in accordance with proper induction or abduction; she thinks that complex explanations (other things being equal) are more likely to be correct than simpler ones, and so on. Nevertheless, she has no idea of the extent of her shortcomings. (Too many of us are often in the position of meaning well with respect to some aim, being nevertheless incompetent at it, and having little or no idea that this is the case.)

To make the comparison between the virtuous and merely well-meaning agents more concrete, consider a merely well-meaning moral and epistemic agent with respect to racial justice. This person values racial equality, but ignores evidence that police statistically treat black and white citizens differently, instead focusing on statistics such as those suggesting that black people are more likely to commit crimes. She has a friend who discusses with her worries that her black son might have a dangerous encounter with police when he is out with his friends at night. Our agent, in trying to console her, says, "Don't worry; as long as he doesn't do anything wrong he'll be fine".¹⁵

Although our agent means well, and values racial justice and knowledge, she doesn't properly value either of them. In believing and acting as

¹⁵See Dotson (2011) for an excellent virtue-theoretic critique of this kind of practice. This is a case where the audience is testimonially incompetent with respect to race (in Dotson's sense). See esp. pp. 246-249 for discussion of a similar example.

she does, she fails to manifest proper respect for what it takes to get onto the facts in this domain. Meaning well just isn't good enough.

Of course, our merely well-meaning epistemic agent is highly unreliable. But her epistemic shortcomings do not stop there. Her perspectival duplicate in the evil demon scenario is intuitively worse off than the virtuous duplicate. This is so despite the fact that from her perspective, she cannot tell the difference between her situation and a virtuous one. Her position is subjectively indiscriminable, and yet she fails to believe rationally. Meaning well does not make it so: merely meaning to believe rationally does not thereby make one believe rationally.¹⁶

If our intuitions about the evil demon scenario are to be taken seriously, we are now confronted with the challenge of explaining the difference in epistemic standing between the virtuous and merely well-meaning duplicates in mental terms, not just in terms of reliability. This issue faces all theories of justification, and it is more difficult to solve than is often acknowledged.¹⁷ For example, consider this passage by Cohen:

Beliefs produced by good reasoning are paradigm cases of justified belief and beliefs arrived at through fallacious or arbitrary reasoning are paradigm cases of unjustified belief. Whether or not reasoning results in false belief, even if this happens more often than not, is irrelevant to the question of whether the reasoning is good. To maintain otherwise would be on a par with confusing truth and validity.

—Cohen (1984), p. 283.

Here Cohen suggests that the question of whether an agent reasons rationally is orthogonal to the question of whether she reasons reliably. But that is not so. Deductive reasoning, after all, is plausibly epistemically valuable precisely because it is conditionally (perfectly) reliable. Thus truth-connectedness, perplexingly, does seem to matter for internalist justification. Moreover and more importantly, the majority of our belief formation and retention does not rely on deductive reasoning, but on the basis of heuristics, induction, and abduction. These kinds of reasoning are not plausibly reduced to logical reasoning; instead, what makes these ways of

¹⁶I here put aside views that entail the opposite conclusion, such as plausibly Foley (1987). I think we can respect the core internalist insight without giving it up.

¹⁷Cohen (1984) is clear that we need a theory of what makes good reasoning good *on the first order*.

forming beliefs rational seems to depend on whether or not they are ways of reliably getting onto the facts.¹⁸

What originally seemed like a clear-cut distinction between externalist kinds of epistemic standing—which have to do with being appropriately hooked up to the world—and internalist kinds of epistemic standing—which have to do with having the appropriate subjective mental life—is starting to look much less clear. How might we articulate what features of the subject’s mental life determine internalist justification without collapsing into a form of externalism, or ending up with the unpalatable consequence that merely meaning to believe rationally makes it so? If we can answer this question, we can solve the new evil demon problem in a truly satisfying way—in a way that does justice to both internalist and externalist insights. In the rest of the paper I will show how my direct virtue epistemology can be extended to do just this.

2 Externalist Justification for Direct Virtue Epistemology

The epistemic theory I defend is a knowledge-first virtue epistemology. It shares with other kinds of reliabilist virtue epistemology the idea that knowledge is an achievement that is due to our epistemic competence, and that epistemic competences are by nature reliable at accomplishing what they are competences to do. However, it is knowledge-first in holding that epistemic competences are competences to *know*, rather than to believe truly, and so the theory is *direct* in the sense that it claims that the competences responsible for knowledge are competences to do that very thing, not to do something that falls short of knowledge. Competences to know must therefore be reliable with respect to knowledge, not just true belief.

Competences to know are reliable but typically fallible; they not only have exercises that are cases of knowledge (manifestations), but they typically also have exercises that are constitutively *failures to know* (degenerate exercises). This feature of competences is central to the view: instead of supposing, as belief-firsters do, that epistemic competences are exercised in a way that is neutral with respect to whether or not they accomplish their aim, I argue that exercises of epistemic competence always entail either success or failure.¹⁹ The success cases (cases of knowledge) are metaphysically and explanatorily more fundamental than the failure cases, however.

¹⁸Williamson (this volume) also makes this point.

¹⁹I hold this to be true for competences more generally. See Miracchi (manuscript) for further discussion.

First, it is essential to epistemic competences that they manifest in cases of knowledge. It is not essential or necessary for them to be able to have degenerate exercises: reliability with respect to knowledge might be perfect, e.g., such as some claim is the case with the *Cogito*.²⁰

Moreover, degenerate exercises are only exercises of competences because their conditions deviate from manifestation conditions; thus degenerate exercises of competence depend on manifestations for their status as epistemic states at all. Being an exercise of competence, rather than being the most epistemically fundamental case, is instead a disjunctive kind—that of either manifesting or degenerately exercising one’s competence—and it is thus metaphysically and explanatorily dependent on manifestations and degenerate exercises.²¹

Nevertheless, the category of *exercise of epistemic competence* does interesting theoretical work. Beliefs have a certain kind of positive epistemic standing in virtue of being members of that category: exercises of competence are as a matter of their nature likely to be cases of knowledge. If knowledge is the fundamental epistemic good—*qua* the achievement of the epistemic domain—and reliability with respect to a good is therefore derivatively a good of that kind, then reliability with respect to knowledge is an epistemic good. A belief is justified in the externalist sense, then, just in case it is an exercise of a competence to know.

According to direct virtue epistemology, not only is justification metaphysically and explanatorily dependent on knowledge, so is belief. Beliefs constitutively aim at knowledge.²² That is, beliefs are just the kind of mental state that aim at knowledge as a matter of their nature. We may now put the point as follows: As the performances that aim at knowledge, beliefs are *the* candidates for knowledge.²³ However, rather than being a unified kind, beliefs admit of importantly different varieties, in accordance with the facts in virtue of which they have knowledge as their aim. In cases of justified belief, it is because the performance is an exercise of competence (a manifestation or a degenerate exercise) that it aims at knowledge, and so

²⁰Sosa (2007), pp. 16-17 also makes the suggestion that the *Cogito* should be thought of as a case of a manifestation of a perfectly reliable competence.

²¹Why do I call exercises of competence a disjunctive *kind*? Isn’t that an oxymoron? As we’ll see, some beliefs can have normative statuses in virtue of being exercises of competence. Thus although the *exercise of competence* is a disjunctive notion, instances of it have properties in virtue of being a member of that kind. Thus *kind*-talk is warranted, at least in my view. Thanks to Neil Mehta for pressing me on this question.

²²Bird (2007) and Sutton (2007) also hold this view.

²³In the sense I am using the term “candidate” here, cases of knowledge are also candidates for knowledge.

is a belief. This is just a special case of the idea that in exercising a competence to A, the agent aims to A.) But there are other ways for a performance to aim at knowing. These are unjustified beliefs.²⁴

We may now put the view about justification slightly differently (though I think equivalently) to what I propose in Miracchi (2015a). According to the theory on offer, an agent's belief that *p* is externalist-justified just in case it is, as a matter of its nature, a *good candidate* for knowledge in the probabilistic sense. Exercises of competence as such (a) aim at knowledge, and so are beliefs, and (b) are likely to be cases of knowledge. Exercise of competence are thus as a matter of their nature good candidates for knowledge in the probabilistic sense.

I now wish to expand this conception of justification and hold that a belief is epistemically justified—in either the externalist or the internalist sense—just in case it is a good candidate for knowledge as a matter of its nature. Moreover, a belief is a good candidate for knowledge if and only if it is an exercise of epistemic competence. Miracchi (2015a) shows how exercises of epistemic competence are good candidates for knowledge in an externalist sense. I will now argue that they are good candidates for knowledge in an internalist sense: all exercises of epistemic competence thereby meet a mental requirement for being knowledge. In the next section I will explain what that mental requirement is, and in section 4 I will explain why all and only exercises of epistemic competence meet that requirement.

3 Rational Believing Is A Kind of Properly Valuing Knowledge

A promising place to start is by looking at some insights from recent work on *derivative value*, and those in epistemology who are already applying it to the epistemic domain. Several people have argued recently, perhaps most notably Thomas Hurka (2001), that some acts and attitudes are valuable because they instantiate or manifest proper ways of valuing something valuable.²⁵ For example, it is not only good to provide food and shelter to the homeless, it is also good to value the acts of providing food and shel-

²⁴See Miracchi (2015a) pp. 50-51 for further discussion.

²⁵This has been of particular interest in the literature on *fitting attitude* theories of value, for those theories try to analyze (certain) values in terms of being worthy of certain attitudes. Regardless of whether this project is on the right track, it has reminded us that certain acts are valuable because they manifest proper ways of valuing the valuable. The articulation of the view in terms of *manifestation* of a proper way of valuing is due to Kurt Sylvan (manuscript), but I think it is a useful way of clarifying explaining Hurka's original view, rather than a development of the view.

ter to the homeless—perhaps by writing a journalistic piece about an organization that does so effectively. Perhaps such writing will increase the number of donations, and so increase the number of homeless people given food and shelter, and so be instrumentally good. However, even if the article were to fail in this regard, it would nevertheless be good merely for the reason that it manifests the author’s valuing the providing of food and shelter to the homeless.

Kurt Sylvan (manuscript) recently pursues this line of thought in providing an account of epistemic value, where truth is the fundamental epistemic value, and cases of believing that properly value the truth are thereby derivatively (non-instrumentally) valuable. He claims that “beliefs are epistemically valuable because they manifest certain ways to place value on accuracy in thought”.²⁶ Sylvan then claims that he can analyze certain epistemic normative properties such as rationality, coherence, and knowledge in terms of different ways of valuing the truth.

While I am less optimistic about being able to account for these epistemic normative properties in the way that Sylvan does, I think he is on the right track in investigating the kind of epistemic standing that the internalist is getting at when she claims that the subject in the evil demon scenario is still justified in believing as she does.²⁷ Properly valuing an epistemic good is clearly something that is inherently first-personal, that has to do with how we mentally, perspectively, proceed in our epistemic inquiries. By placing attention on whether or not the subject properly values the truth (or knowledge!) in believing as she does, we are placing our attention on something that is clearly a feature of her mental life.

As this stands however, it won’t quite do, for two reasons. First, I need to explain why the kind of proper valuing I am claiming is constitutive of epistemic rationality is plausibly something that all beliefs have, and is not overly intellectualized.²⁸ Second, it is important to distinguish the kind of derivative value that beliefs can have from the kind that performances which are not beliefs can have. For example, one way of valuing knowledge is to create schools. But the act of creating schools, if epistemically valuable, is valuable in a very different sense from the epistemic value of believing rationally. It is certainly not epistemically *rational* in the same sense that beliefs are epistemically rational. Hurka’s and Sylvan’s accounts,

²⁶Sylvan (manuscript) p. 4.

²⁷Sylvan is clear that he does not mean for such an account to be a contender to a moderate reliabilism, which claims that reliability is an epistemic good; rather he aims to be augmenting such a view—accounting for a more internalist kind of positive epistemic standing.

²⁸Hurka and others face an analogous challenge with respect to the moral domain.

however, do not make this distinction.²⁹

Note that, in this case, analogously to the journalism case above, the claim is not that writing journal articles or opening schools is epistemically valuable because it is a way of promoting knowledge, but that it is epistemically valuable because it embodies the proper valuing of knowledge. This kind of proper valuing, however, is different and more removed from the kind of proper valuing our beliefs have because we value what it takes to know in believing as we do. How shall we understand the difference?

4 Competent Perspectives

Virtue epistemology, and in particular my direct virtue epistemology, can help with both of these issues. First I will take up the question of distinguishing the kind of derivative value beliefs can have from other kinds of derivative value. Then, I will answer the charge of over-intellectualizing rationality.

According to virtue epistemology, epistemology is a *performance domain*. This means its normativity is structured in terms of certain aims that are fundamental to the domain, and the agency involved in attaining those aims. The primary bearers of epistemic value are the performances that are candidates for being attainments of the fundamental epistemic aim(s) of the domain. As discussed in the previous section, for direct virtue epistemology knowledge is the fundamental aim, and beliefs are the performances that are candidates for knowledge. This is why beliefs are the immediate bearers of epistemic properties.³⁰

Other performances may bear epistemic properties only as they relate to the performances that are candidates for knowledge. Opening a school is an example of a performance that is not a candidate for knowledge, and therefore has epistemic status only at a remove. Opening a school does not aim at knowledge in virtue of its nature. Only beliefs do that.³¹ One who opens a school with the right motives both increases the amount of knowledge in the world and manifests proper valuing of knowledge, but not by performing in a way that is itself a candidate for knowledge. Thus, although it is both instrumentally valuable and manifests proper valuing

²⁹Sylvan does not discuss this problem for his view.

³⁰Epistemic agents too are immediate bearers of epistemic properties, because they are the ones who achieve knowledge.

³¹This avoids Berker (2013a,b) style worries. When beliefs are formed or maintained in the aim of creating further knowledge, they are not aiming at knowledge *qua* candidates for knowledge.

of knowledge, it is epistemically valuable only in a derivative sense.

Within the virtue-theoretic framework I have offered here, we can now provide a motivated restriction on the kind of proper valuing of knowledge that is constitutive of epistemic rationality: it is properly valuing knowledge in aiming to know. In other words, the kind of proper valuing we are after is a kind of practical valuing: *A belief is rational just in case the epistemic agent properly values what it takes to achieve her aim of knowledge in believing as she does.*

Now we can address the other worry for the account, namely that it over-intellectualizes epistemic rationality. Do we really properly value knowledge as the aim of our performance every time that we know? Of course, sometimes we know things we would rather not know. We might even wish that we could allow other more practical considerations to sway us. However, I think that even in such cases there is a sense in which we properly value knowledge as the aim of our doxastic performance.

The epistemically virtuous agent does not experience a blind attraction to believing in a way that is, as a matter of how things turn out, knowledge; rather, the fact that this is so guides her in her reasoning. In manifesting her competence, the virtuous epistemic agent is attracted to certain patterns of reasoning precisely because they are ways of acquiring and maintaining knowledge. It is precisely because the evidence unequivocally points to p that one believes p , even when one would rather not do so. The “because” as I am using it here is not merely causal. It entails a kind of sensitivity, from the subject’s own perspective, to the fact that to perform in a certain way is to provide oneself with, or maintain, knowledge.

This is exactly the kind of feature we have been looking for: believing and reasoning in certain ways *because* they are ways of knowing suffices for the agent to properly value knowledge as the aim of belief in believing as she does. After all, to perform in a certain way because doing so *would* be an achievement of one’s aim is plausibly the best way to value what it takes to achieve one’s aim in performing as one does.

To say that the virtuous epistemic agent is attracted to certain ways of reasoning because they are ways of acquiring and maintaining knowledge does not commit one to the claim that the agent *believes*, or can articulate, this attraction.³² Nor does it commit one to the claim that the agent has a *desire*, or other pro-attitude, to believe in a way that is a way of knowing distinct from her coming to believe or her maintaining her belief. This

³²Note also that if we required knowledge, this would immediately lead to a vicious regress.

would over-intellectualize rationality, as has been widely noted.³³ Rather, I am describing *what it is* to come to believe or maintain one's belief in a competent way on the first order. It has an ineliminably perspectival aspect to it.³⁴

Although this was not the focus of Miracchi (2015a), the sense in which the subject has an aim to know in exercising her epistemic competence was always supposed to be perspectival. As opposed to some other virtue epistemologists such as Ernest Sosa who assimilate the aiming of everyday belief to biological functioning, for me it is very important that the aim that is constitutive of belief is mental in a way that the function of a heart to pump blood is not.³⁵ Rather than thinking of aims to know as biological or evolutionary, we should think of them as a distinctively mental kind of directedness.

If we do not, then it is by no means clear that we are dealing with performances of *the epistemic agent* in any important sense. Virtue epistemology, which was designed to center the agent in our epistemic theorizing, thus falls prey to the same problems that reliabilism does: it turning us (qua epistemic beings at least) into mere machines that are reliably hooked up to the world. We are clearly much more than that, and it is the task of the naturalistically-minded epistemologist to articulate how this might be the case.³⁶

Once we have gotten this far, however, the idea that the way in which we aim at knowing when we manifest our epistemic competence entails proper valuing of that aim is not far off. In such cases, having the aim of knowledge in believing as one does *is* competently aiming for knowledge. It is being drawn to take certain objections seriously, to revise one's commitment in light of (seeming) counter-evidence or to undermine that counter-evidence, and so on. It is being drawn to certain ways of believing because they are ways of knowing, and that is just what it is to properly

³³E.g. see Cohen (1984)'s criticisms of Lehrer. Sylvan (manuscript) is also clear that understanding rationality to be proper valuing does not commit one to such a view.

³⁴It is plausible that an analogous phenomenon occurs in the moral domain. The virtuous moral agent, in exercising her competence, just is just motivated to do what is right. She is drawn, first-personally, perspectively, towards an action and in doing so she manifests properly valuing doing what is right.

³⁵See esp. Sosa (2015), p. 20. I think there are many problems with a biological or evolutionary account of epistemic aims, but this is not the place to discuss them.

³⁶To be clear, the claim is not that we need a non-naturalist or anti-physicalist account to explain these mental properties. However, acknowledging these mental properties and their centrality to epistemology is crucial for developing an adequate theory.

value knowledge as the aim of one's performance.³⁷ This should allay our fears that the account over-intellectualizes rationality. Rather, it appeals to the most basic perspectival-motivational features that are present in any manifestations of competence that are properly called performances by the agent.

However, I haven't yet explained what happens in a case of rationality that falls short of knowledge. One can only believe in a certain way because it is a way of knowing if that way of believing is indeed a way of knowing. What, then, about rational beliefs that fall short of knowledge? In such cases, we typically imagine a subject situated so that her belief falls short of knowledge, and yet her case is indiscriminable from one where she knows; it is for her *as if* she were acquiring or maintaining knowledge, even though she is not. But again we must remember that it can't be indiscriminability as such that does the trick. Recall our merely well-meaning agent. Although she believes in a way that is indiscriminable from a way of knowing, she doesn't believe rationally. We have to explain what is going on with the merely rational believer in a way that goes beyond reference to the indiscriminability of her situation.

Here direct virtue epistemology is again well-poised to provide an answer. As noted above, according to most kinds of virtue epistemology, epistemic competences are typically fallible. Certain conditions might preclude agents from achieving their aims, while nevertheless they perform competently. In such cases the exercise, though degenerate, is still fully *competent*. The agent was merely unlucky.³⁸

Moreover, it is plausible that the reason why the rational agent's situation is indiscriminable from a case of knowledge is that she exercised her competence.³⁹ We can think of what happens here as analogous with perceptual illusions: in cases of illusion, a thing seems to be a certain way even though it is not. What explains why things look that way are the very

³⁷As an example, consider an epistemic case, where an agent competently deduces q from p and *If p then q* , thereby coming to know that q . When the agent is motivated to infer q , she is properly aiming at knowledge, aiming at knowledge in a way that properly values what it takes to know. This agent may not have a concept of *modus ponens*. She may not be able to tell you that the form of inference is valid. However, as long as it is this property of the inference that is perspectivally guiding her in believing as she does, she properly values what it takes to know.

³⁸Of course it behooves an agent to continually try to hone her competences to reduce the probability that she will be undone by bad luck, but that does not mean that in such cases the fault for her failure lies with her.

³⁹Presumably, a fault *would* be with you if you had a sense that you were failing to get onto the facts.

same competences that explain why things look the way they are to you in the good case. Much work in vision science presupposes that this is true. This is why illusions are so empirically interesting—they reveal something about how the competences responsible for veridical perception work.⁴⁰

Similarly, we might think of the rational bad case as a case of epistemic illusion: a certain way of believing seems to be a way of knowing, and this is explained by appeal to the same competence that provides you with knowledge on other occasions. So, although you don't believe as you do because so believing is a way of knowing (you can't, because it's not), you still competently, perspectively, aim at knowledge. In cases of failure, the virtuous epistemic agent competently, albeit mistakenly, believes as she does because she competently takes so believing to be a way of knowing. Her failure does not reflect badly on how she was proceeding; she was merely unlucky.⁴¹ As such, she still properly values knowledge as the aim of her performance.

We now have what are plausibly necessary and sufficient conditions for properly valuing knowledge in believing as one does, and so for rational belief: *in cases of rational belief one believes in a certain way because one competently takes that way to be a way of knowing. In so doing, one properly values knowledge as the aim of one's belief.* Just as in the case of externalist justification, the mental performance that satisfies this requirement differs in the cases of knowledge and rational false or Gettiered belief. When the exercise of competence is a manifestation, the agent believes in a certain way because it is a way of knowing. When the exercise of competence is degenerate, the agent believes in a certain way because she competently (albeit mistakenly) takes a way of believing to be a way of knowing. In both cases, the agent properly practically values knowledge as the aim of her performance because she competently takes the way in which she believes to be a way of knowing.

This account of proper valuing knowledge in believing as one does allows us to distinguish the virtuous epistemic agent from the merely well-meaning and virtuous agents in mental terms. The merely well-meaning agent is not epistemically competent: she is not disposed to believe in certain ways because those are ways of knowing. As such, even though she cannot tell that she is failing to know, she does not competently aim, and so does not properly value her aim in the sense at issue.⁴²

⁴⁰See e.g. Palmer (1999).

⁴¹This discussion is closely related to interesting issues of *direction of fit* which I cannot get into here.

⁴²Of course, sometimes one might be faultless for failing to have a competence—perhaps

The view I am advocating does claim that there is a kind of faultlessness in cases of rational false or Gettiered belief, and so one might wonder whether it falls prey to the objections I made against appeal to excuses in the introduction. However, this new approach is importantly different. First, it explains why certain cases are faultless by appeal to what is going *well* in the situation, namely that the agent exercises her competence. Because the agent exercises her epistemic competence (albeit degenerately), she properly values knowledge as the aim of her performance in believing as she does. Even in the degenerate case her belief is by nature a good candidate for knowledge because it manifests proper valuing of knowledge as the aim of her belief.

The account also explains why indiscriminability from cases of knowledge seems to matter, and why (only) certain cases of mere indiscriminability are faultless. When the indiscriminability is due to an epistemic competence being exercised, the agent competently takes a certain way of believing to be a way of knowing. These are the only cases of indiscriminability that are cases of rational belief. We can thus distinguish the merely well-meaning epistemic agent from the truly rational one.

Lastly, the account of rationality on offer here presents it as importantly the same in kind as externalist justification. Exercises of epistemic competence are *the* good candidates for knowledge, both in an externalist and an internalist sense. When an agent exercises her epistemic competence, she believes in a way that is thereby likely to be knowledge (externalist sense), and she also believes in a way that properly values knowledge as her aim (internalist sense). Moreover, these two features are not independent of one another. It is no accident that the beliefs that are externalist-justified are also internalist-justified, because the very facts that are constitutive of competence possession determine both the reliability of one's exercises and the features of one's epistemic perspective. As such, it avoids the problems set out in the introduction for theories on which rational false or Gettiered belief is excusable failure.

5 Conclusion: Diagnosing the New Evil Demon Problem

So far, I have explained the difference between epistemically virtuous, vicious, and merely well-meaning agents in the actual world. But how does

because of a developmental situation. But we can distinguish these two kinds of faults just like we can distinguish two kinds of luck: there's (bad) luck you have for failing to have certain competences, and then there's (bad) luck you have for failing to manifest the competences you do have. Only the latter are directly relevant to epistemic standing.

this relate to our original question, namely solving the new evil demon problem? The perspectival duplicates are brains in vats, and so don't have any epistemic competences that satisfy the requirements of direct virtue epistemology (they are completely unreliable).

However, recall section 1. There I claimed that in considering the evil demon scenario, we are conceiving of a case where the subject's epistemic perspective remains the same even though her reliable connection to the world is severed. This teaches us that epistemic agents have a certain kind of epistemic standing in virtue of their mental lives, and not directly in virtue of their reliable connections to the world. A solution to the new evil demon problem requires explaining the epistemic differences between cases of rational and irrational belief in terms of the subject's mental perspective. If we have good reason to think that the evil demon scenario is metaphysically impossible, we do not need to appeal to features of the subjects' mental lives that (metaphysically) could obtain in the evil demon scenario.

As discussed in that section, I doubt that any account of internalist justification will be able to explain the difference between the virtuous and merely well-meaning duplicates in a way that avoids appeal to mental features that cannot be had in the evil demon scenario. This is for the simple reason noted above, that what counts as good reasoning (except for logical cases) depends on what it takes to get onto the facts in the agent's world. Inductive, abductive, and heuristic kinds of reasoning are all good or bad in large part because of how the world outside one's head is.

On a knowledge-first virtue-theoretic approach, this is to be expected. Competences to know are fundamental to epistemic evaluation. We cannot explain what makes a belief rational except in terms of what is required for knowledge. Although properly valuing the valuable is a distinctively first-personal phenomenon, then, and part of one's mental life, it is not independent from one's relation to the facts. There are mental differences between the virtuous, vicious, and merely well-meaning duplicates, but in order to explain them we need to appeal to how these agents are situated in their worlds.

If this is correct, then we may reject the metaphysical possibility of the evil demon scenario in a motivated way. The mental feature constitutive of epistemic rationality—proper valuing of knowledge as the aim of one's performance—could not be had in the evil demon scenario. However, as long as we are allowing that the evil demon scenario has perspectival duplicates, we allow ourselves the metaphysical impossibility that the subject believes in a way that properly values knowledge, even though she is not

reliably hooked up to the world. It is important, then, that the evil demon scenario is metaphysically impossible, but not in the way we originally might have supposed. We do not get to reject the case as irrelevant to our epistemic theorizing. We do, however, allow ourselves to appeal to the agent's connection to the world when explaining the differences between various first-personal perspectives.

By accepting the world-dependence of our perspectival lives, we arrive at a unified virtue epistemology. We can explain what justification and rationality have to do with one another, and with knowledge. We can also answer a question that has long been plaguing internalists, namely why certain well-meaning agents are rational, and others not. This problem seemed intractable, I will now suggest, because we were supposing epistemology to be independent of philosophy of mind. We were assuming that we could put aside discussion of how the world contributes to our perspective on it in discussing epistemic standing, but we cannot do so. Our connection to the world does not merely reliably hook us up to it, so that we "produce" beliefs that are likely to be true; it provides us with the kind of grip on the world that can guide us, from our own perspectives, towards knowledge.

References

- Berker, S. (2013a). Epistemic teleology and the separateness of propositions. *Philosophical Review*, 122(3), 337–393.
- Berker, S. (2013b). The rejection of epistemic consequentialism. *Philosophical Issues*, 23, 363–387.
- Bird, A. (2007). Justified judging. *Philosophy and Phenomenological Research*, 74(1), 81–110.
- Brueckner, A. (1986). Brains in a vat. *Journal of Philosophy*, 83(3), 148–167.
- Cohen, S. (1984). Justification and truth. *Philosophical Studies*, 46(3), 279–295.
- Comesaña, J. (2002). The diagonal and the demon. *Philosophical Studies*, 100, 249–266.
- Dotson, K. (2011). Tracking epistemic violence, tracking practices of silencing. *Hypatia*, 26(2), 236–257.

- Foley, R. (1987). *The Theory of Epistemic Rationality*. Harvard University Press.
- Hurka, T. (2001). *Virtue, Vice, and Value*. Oxford University Press.
- Ichikawa, J. J. & Jenkins, C. I. (manuscript). On putting knowledge 'first'.
- Lehrer, K. & Cohen, S. (1983). Justification, truth, and coherence. *Synthese*, 55(2), 191–207.
- Littlejohn, C. (2012). *Justification and the Truth Connection*. Cambridge University Press.
- Miracchi, L. (2015a). Competence to know. *Philosophical Studies*, 172(1), 29–56.
- Miracchi, L. (2015b). Knowledge is all you need. *Philosophical Issues*, 25(1), 353–378.
- Miracchi, L. (manuscript-). Achievements and exercises: A theory of competence.
- Palmer (1999). *Vision Science: From Photons to Phenomenology*. MIT Press.
- Putnam, H. (1981). Brains in a vat. In *Reason, Truth, and History*. Cambridge University Press.
- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, volume 1. Oxford University Press.
- Sosa, E. (2015). *Judgment and Agency*. Oxford University Press.
- Sutton, J. (2007). *Without Justification*. Cambridge, MA: MIT/Bradford.
- Sylvan, K. (manuscript). What isn't the truth-connection?
- Williamson, T. (this volume). Justifications, excuses, and skeptical scenarios. In F. Dorsch & J. Dutant (Eds.), *The New Evil Demon: New Essays on Knowledge, Justification and Rationality*. Oxford University Press.