

UC Irvine

UC Irvine Previously Published Works

Title

Moderate Knowledge Externalism

Permalink

<https://escholarship.org/uc/item/4pr584sm>

Author

Pritchard, Duncan

Publication Date

2023-08-01

DOI

10.1093/oso/9780198866749.003.0006

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

MODERATE KNOWLEDGE EXTERNALISM

DUNCAN PRITCHARD

University of California, Irvine

dhpritch@uci.edu

ABSTRACT. A case is made for a moderate version of form of epistemic externalism about knowledge. It is argued that the general structure of knowledge is along the lines set out by *anti-risk virtue epistemology*. Since such a proposal makes no essential demand that one's cognitive success be significantly supported by reflectively accessible rational support, epistemic externalism about knowledge results. Nonetheless, anti-risk virtue epistemology offers an account of knowledge that is entirely built around a conception of cognitive responsibility. This ensures that it is able to accommodate supposedly internalist intuitions about the relationship between knowledge and cognitive responsibility. Moreover, it is argued that this proposal is entirely consistent with the idea that the knowledge possessed by cognitively developed agents characteristically involves a significant level of reflectively accessible rational support. Indeed, it is claimed that not only does the knowledge possessed by cognitively developed agents enjoy such rational support but also, in line with *epistemological disjunctivism*, that such reflectively accessible rational support will also often include *factive* reasons. It is argued that far from being in tension with epistemic externalism about knowledge, epistemological disjunctivism offers a compelling way for a moderate epistemic externalism to account for the role of reasons within knowledge. The resulting account of knowledge, while essentially a form of epistemic externalism, can thus nonetheless incorporate many of the considerations that drive the adoption of epistemic internalism about knowledge.

1. MAPPING THE TERRAIN

Let's begin by defining terms. Call an *epistemic condition* any necessary condition for knowledge, in addition to true belief (where it is left open how many epistemic conditions there are). I will take *epistemic internalism about knowledge* to consist in the claim that a necessary condition for knowledge is that one's belief in the target proposition satisfies an internalist epistemic condition. *Epistemic externalism about knowledge* will thus consist in the denial of this claim, and thus amounts to the

thesis that one can have knowledge without satisfying an internalist epistemic condition. I will understand an *internalist epistemic condition*, in turn, as a condition on knowledge which demands that the subject has reflectively accessible reasons for believing the target proposition to be true.¹ An *externalist epistemic condition* will thus be any condition on knowledge which does not involve this demand.

This way of thinking about the epistemic internalist/externalist distinction and its application to the theory of knowledge is, of course, at least partly stipulative, but I think it nonetheless broadly captures some important distinctions in post-Gettier epistemology. In terms of this framework, for example, we can consider the classical tripartite account of knowledge as a particularly pure form of epistemic internalism about knowledge, in that satisfying an internalist epistemic condition (justification, classically understood) is not only necessary for knowledge but also, with true belief, sufficient (i.e., there is only epistemic condition on knowledge, and it's an internalist epistemic condition).² Post-Gettier, however, the interesting battleground regarding knowledge is between impure forms of epistemic internalism (i.e., which have a necessary internalist epistemic condition, but which also include external epistemic conditions) and epistemic externalist proposals which reject the need for an internalist epistemic condition altogether (such as process reliabilism as applied to knowledge).³

With the foregoing description of the topography of the epistemic externalism/internalism debate in mind, I will be arguing for three main theses. First, I will be articulating an externalist account of knowledge in general that I call *anti-risk virtue epistemology*. In particular, I will be explaining how this view can offer us a robust way of understanding the manner in which knowledge requires cognitive responsibility without thereby demanding that the subject's belief satisfies an internalist epistemic condition.

Second, I will be arguing that even if one embraces anti-risk virtue epistemology—and thus epistemic externalism about knowledge in general—that is entirely consistent with holding that the knowledge possessed by cognitively developed agents often satisfies an internalist epistemic condition. In particular, I will be contending that paradigmatic forms of perceptual knowledge in fact satisfy a particularly robust rendering of the internalist epistemic condition, by involving a supporting reason that is *factive* (i.e., which entails the believed proposition that it is a reason for). This position is known as *epistemological disjunctivism*.

Third, I will maintain that not only is epistemological disjunctivism compatible with an externalist account of knowledge like anti-risk virtue epistemology, but that the role it offers for factive reasons within a theory of knowledge is crucial to making sense of how rational support relates to knowledge within an externalist account. The alternative, as I explain, is a conception of

rational support that runs in parallel to the kind of cognitive responsibility captured by epistemic externalism, but which is ultimately disengaged from it. What we need to make sense of is how, in the right epistemic conditions, the rational support we have for our beliefs is no less hooked up to the world than the externalist epistemic support that these beliefs enjoy.

2. TWO KINDS OF COGNITIVE RESPONSIBILITY

The attraction of epistemic internalism about knowledge lies in the notion of cognitive responsibility. If one's knowledge doesn't satisfy an epistemic internalist epistemic condition, such that one has good, reflectively accessible, reasons in support of one's belief, then in what sense is one's knowledge really attributable to one's cognitive agency at all? This point becomes especially clear if one focusses on an austere form of epistemic externalism about knowledge, such as process reliabilism. If reliability is a route to knowledge, regardless of whether one has reasons available to one in support of the target proposition—including regarding whether one's belief is formed via a reliable process—then in what sense does this knowledge have anything to do with one's exercise of cognitive agency? Indeed, couldn't it just as well be a matter of brute luck that one knows, at least from one's own perspective, in that one has simply chanced upon a true belief via a reliable process?⁴

In contrast, if one's true belief is grounded in good reflectively accessible reasons, then it plausibly is now attributable to one's cognitive agency, as one is able to offer an appropriate account of why one regards the proposition as being true. Note that given we are working, post-Gettier, with at most impure epistemic internalism about knowledge, then there is no pretense that such reasons would entail the truth of the target proposition, or otherwise would suffice (in conjunction with true belief) for knowledge by themselves. One is thus still to some extent hostage to the kind of external factors obtaining that are at issue in the relevant external epistemic conditions in play. Nonetheless, so the thought goes, what's important for knowledge is just that one can tell such a story about cognitive agency, and that epistemic internalism about knowledge is uniquely placed to capture this insight.

I don't want to deny the importance to our epistemic framework of reflectively accessible reasons. Indeed, as already indicated, I think that often we are in possession of such reasons of a very special kind, in that they actually entail the truth of the target proposition. But my current interest is in the idea that epistemic internalism about knowledge is the only game in town when it comes to accounting for the importance of such reasons. In particular, I want to suggest that the notion of cognitive responsibility in play here is in fact much broader than the epistemic internalist

would have us believe.

That there is such a broader notion of cognitive responsibility is an insight that we gain from virtue-theoretic accounts of knowledge. The contrast between process reliabilism and virtue reliabilism is a case in point. We can take the latter to be essentially a refinement of the former. One of the key problems that afflicts process reliabilism is that it treats any reliable belief-forming process as knowledge-conducive, and yet some reliable belief-forming processes appear to have nothing whatsoever to do with knowledge. Consider, for example, Alvin Plantinga's (1993a, 195-98) famous case of the brain lesion that causes the agent concerned (who is oblivious as to the causal processes in play) to reliably form a true belief that she has a brain lesion. Getting to the truth in this way looks very much like dumb luck. The crux of the matter is that the reliability of this belief-forming process seems irrelevant to knowledge because it has nothing to do with the subject's cognitive agency. Indeed, this scenario concerns a *malfunction* of one's cognitive agency, in that it is *despite* one's cognitive agency that one ends up forming a true belief via a reliable process.

Virtue reliabilism aims to avoid the problems posed by cases like this by restricting the class of reliable belief-forming processes that are relevant to knowledge. In particular, it is not true belief formed via any reliable belief-forming process that puts one into the market for knowledge, but only those skill-like reliable belief-forming processes that are integrated within one's cognitive character, and which are thus manifestations of one's cognitive agency. We can thus explain why one's reliable faculties and cognitive skills can be knowledge-conducive while brain lesions and such like, even when they generate a reliable belief-forming process, are not knowledge-conducive.⁵

The general point in play here is that knowledge demands cognitive responsibility at least in the following sense: that one's cognitive success (i.e., true belief) is significantly creditable to one's cognitive agency. This means that one's cognitive agency is playing a significant role in the causal explanation of one's cognitive success. Call this *veritic cognitive responsibility*. Merely forming one's true belief in a reliable way doesn't suffice in this regard, as one's cognitive success might have nothing to do with one's manifestation of cognitive agency, as when the reliability in question arises out of a cognitive malfunction rather than a cognitive skill. In contrast, when one gets to the truth by employing one's reliable cognitive traits that are integrated into one's cognitive character—and thus count as one's cognitive skills—then one's cognitive success is significantly attributable to one's cognitive agency.

The reason why veritic cognitive responsibility is important for our purposes is that we are capturing a notion of cognitive responsibility which doesn't essentially trade in reflectively accessible reasons. Imagine for a moment a subject whose perceptual faculties are functioning appropriately in suitable cognitive conditions, and hence is reliably forming true beliefs as a result,

but who lacks reflectively accessible grounds in support of her beliefs. What's interesting about such a case is that it looks very different to someone merely coming across the truth through dumb luck. The point is that getting to the truth via dumb luck has nothing to do with one's manifestation of cognitive agency. In contrast, when one appropriately employs one's cognitive skills and gets to the truth as a result, then one's cognitive success *is* significantly attributable to one's cognitive agency, even if one lacks reflectively accessible grounds in support of the target belief. It follows that veritic cognitive responsibility captures a general kind of cognitive responsibility that doesn't essentially appeal to the possession of reflectively accessible reasons.

Notice that in stating the foregoing one need not be claiming that the possession of reflectively accessible reasons is unimportant to knowledge. The point is just that one can have knowledge, and hence exhibit a sufficient level of cognitive responsibility for knowledge, without being in possession of such reasons. In fact, when it comes to cognitively developed agents, I think it is clear that our beliefs tend to go hand-in-hand with the possession of reflectively accessible reasons. Indeed, to stand in the 'space of reasons' in this way is a core part of what it is to be a cognitively developed agent. But while the possession reflectively accessible reasons might normally be the case, it's not a requirement of knowledge, even though it is a requirement of knowledge that it demands veritic cognitive responsibility.

Moreover, notice that while the possession of reflectively accessible reasons is unnecessary for the kind of veritic cognitive responsibility that knowledge demands, it doesn't follow that the possession of such reasons has no bearing on one's cognitive responsibility. Quite the contrary: one will usually manifest higher levels of cognitive responsibility by in addition possessing such reasons. Consider our agent forming her true beliefs through the employment of her reliable cognitive abilities, while lacking reflectively accessible reasons in support of what she believes. Wouldn't her cognitive success be attributable to her cognitive agency to a higher degree if she also comes to possess good reflectively accessible reasons that align with her manifestation of cognitive ability, such as by becoming aware of her track-record of success in this regard? The point is that in denying the necessity of reflectively accessible reasons for the kind of veritic cognitive responsibility required for knowledge one is not thereby denying that such reasons can often play an important role in determining whether, and to what extent, one exhibits such cognitive responsibility.

It follows that even an epistemic externalist about knowledge ought to not only be willing to grant that the knowledge possessed by cognitively developed agents involves the satisfaction of an internalist epistemic condition, but will also prize the possession of reflectively accessible reasons, as it can ensure a higher level of cognitive responsibility. What makes one an epistemic externalist is just that one doesn't hold that satisfying an internalist epistemic condition is required

for knowledge. In particular, one's true belief can exhibit the kind of veritic cognitive responsibility required for knowledge without satisfying such an epistemic condition.

3. ANTI-RISK VIRTUE EPISTEMOLOGY

In the last section I argued that knowledge demands veritic cognitive responsibility. In this section I want to argue that knowledge doesn't just demand veritic cognitive responsibility, but also a stronger kind of cognitive responsibility, albeit where this stronger notion still falls short of demanding reflectively accessible grounds (and hence doesn't lead to epistemic internalism about knowledge either). The stronger thesis that I have in mind is that one's *safe* cognitive success is significantly attributable to one's cognitive agency. Call this *safe cognitive responsibility*. Moreover, safe cognitive responsibility amounts to a complete account of knowledge, one that is thus both epistemic externalist and centrally concerned with cognitive responsibility. I call this theory of knowledge *anti-risk virtue epistemology*.⁶

Safety demands that one's cognitive success, so formed, could not have easily been a cognitive failure (i.e., false belief).⁷ Put another way, there is not a close possible world where one forms one's belief on the same basis as in the actual world and one's belief is false. In earlier work I have motivated this constraint on knowledge in terms of a methodology that I called *anti-luck epistemology*.⁸ Roughly, one first offers an account of luck, then one offers an account of the specific sense in which knowledge is incompatible with luck, and then one puts the two parts together to articulate an anti-luck condition on knowledge. To this end, I offer a modal account of luck and a specification of the veritic epistemic luck that's inconsistent with knowledge, and this leads directly to the safety condition. We thus have an independent motivation for this constraint on knowledge.⁹ These days I make the same point by appealing to risk rather than luck, and hence we get *anti-risk epistemology*. The details behind the switch aren't important for our purposes; what is important is rather the guiding idea that knowledge is incompatible with a high degree of epistemic risk, and that this is captured by safety.¹⁰

That knowledge demands safety is widely accepted in the literature, and we've already witnessed the rationale for thinking that knowledge demands veritic cognitive responsibility.¹¹ But why would it be important to knowledge that it in addition demands safe cognitive responsibility? We can understand why by considering two kinds of cases: (i) veritic cognitive responsibility involving an unsafe true belief; and (ii) veritic cognitive responsibility involving a safe true belief, but where one's safe cognitive success is not significantly attributable to one's cognitive agency. We will take these cases in turn.

In order to understand what is at issue in (i), we first need to introduce a distinction between two ways in which a success can be unsafe, and thus high risk: *intervening risk* and *environmental risk*.¹² Consider an archer who skillfully fires an arrow at a target, and who hits the target. The twist in the tale, however, is that a dog playfully jumped up and caught the arrow mid-flight and placed it in the target. We thus have skillful agency on display and also a successful outcome, but where the success could very easily have been a failure (e.g., had the dog run off with the arrow instead), and hence is unsafe. This is a case of intervening risk, in that something intervenes between the skillful attempt and the successful outcome, such that there was a high risk of failure. In cases of intervening risk, the success is not significantly attributable to the subject's skillful agency, but is rather attributable to the intervening factor, in this case the contribution of the playful dog.

But consider now a corresponding case of environmental risk. We run the case just as before, except that now instead of the dog playfully grabbing the arrow in mid-flight, he instead narrowly fails to catch it. The arrow therefore passes the dog and goes on to hit the target. We thus have skillful agency on display and also a successful outcome. While there is no intervention in this case, the subject's success is still unsafe, given how the dog could very easily have intervened; there was still a high risk of failure. This lack of safety is, however, due to environmental factors rather than because of anything intervening between the skillful attempt and the successful outcome. Interestingly, unlike cases of intervening risk, the success in this case *is* significantly attributable to the subject's skillful agency, even though it is unsafe. The point is that one's success can be modally fragile and yet no less attributable to one's agency, at least so long as nothing in fact intervenes between the skillful attempt and the successful outcome.¹³

This contrast between intervening and environmental risk manifests itself in epistemic cases. Intervening epistemic risk will be familiar from standard Gettier-style scenarios, as these usually involve the combination of cognitive skill and cognitive success, but where something intervenes to ensure that the latter is not attributable to the former. (Consider, for example, the famous 'sheep' case, where the farmer sees what he thinks is a sheep in the field, and so truly believes that there is a sheep in the field, but he is in fact looking at is a big hairy dog, which is obscuring from view the genuine sheep behind).¹⁴ In contrast, environmental epistemic risk is found in the barn façade case and other non-standard Gettier-style cases that have a parallel structure.¹⁵ Recall that the subject genuinely sees a barn in this scenario, but that her cognitive success is nonetheless unsafe because she happens to be in an environment in which what looks like a barn is normally a barn façade.

We get the same contrast between intervening risk and environmental risk in the epistemic cases as in the non-epistemic examples just considered. In the former, the intervening epistemic

risk ensures that the subject's cognitive success is no longer significantly attributable to her cognitive agency. In contrast, in cases of environmental epistemic risk the subject's cognitive success is significantly attributable to her cognitive agency, even though it is unsafe. Neither case amounts to knowledge, even the case of environmental epistemic risk where the subject exhibits veritic cognitive responsibility.¹⁶ The natural explanation for this is that the cognitive success is unsafe (and thus epistemically risky) in both cases, as the subject's belief-forming process could have easily resulted in a false belief. Knowledge is, however, incompatible with high levels of epistemic risk, which is why it demands safety. We can thus see why veritic cognitive responsibility involving an unsafe true belief, as in cases of environmental epistemic risk, doesn't amount to knowledge.

Note, however, that this conclusion needs to be properly understood. In particular, as it stands it might be thought that the way to resolve the issue would be to demand that knowledge requires veritic cognitive responsibility and, in addition, safety. But this would be a mistake. In particular, while knowledge does demand both veritic cognitive responsibility and safety, these are not to be understood as distinct constraints on knowledge.

We can see this point by considering cases that fall under (ii), such that there is veritic cognitive responsibility involving a safe true belief, but where the subject's safe cognitive success is not significantly attributable to her cognitive agency. Consider the following variation on the barn façade case. In this scenario we keep everything fixed from before except that we add a helpful demon who is determined to ensure that the subject forms a true belief in this regard. In particular, while the subject in fact looked at a genuine barn, had she instead opted to look at one of the fake barns in the vicinity, then the helpful demon would have instantly switched it for a real barn to ensure that the subject's belief remains true. The environmental epistemic risk is thus cancelled out by the helpful demon, and the subject's belief is hence nonetheless safe (even though the helpful demon does not in fact do anything).¹⁷

As with the original barn façade case, the subject's cognitive success is significantly attributable to her cognitive agency (and hence it manifests veritic cognitive responsibility), as it remains true that nothing has intervened between the skillful formation of belief and cognitive success. Moreover, the belief is in addition safe, thanks to the watchful eye of the helpful demon, as given how the belief was formed this is not a cognitive success which could have easily been a cognitive failure. And yet the belief is no more a candidate for knowledge than the original barn façade belief. Furthermore, the natural diagnosis for why this is so is that the safety of the cognitive success has nothing at all to do with the subject's manifestation of cognitive agency, but is rather solely the result of the presence of the helpful demon. In particular, what is lacking in this case, unlike in genuine cases of knowledge, is that the subject isn't exhibiting safe cognitive

responsibility, but merely the combination of safety and veritic cognitive responsibility.

Understanding knowledge in terms of safe cognitive responsibility, as anti-risk virtue epistemology proposes, offers us a complete externalist theory of knowledge, in that it captures all cases. It would obviously take us too far afield to defend this claim in full here, but let me at least give the reader a flavor of the considerations in its favour by comparing it to a recent account of knowledge that is broadly cast along the same lines (and which is also a form of epistemic externalism).¹⁸ According to this proposal, which I have elsewhere called *robust virtue epistemology*, we should understand knowledge as demanding a strong version of veritic cognitive responsibility, whereby the agent's cognitive success should be *primarily*, rather than merely significantly, attributable to the agent's cognitive ability.¹⁹ In particular, the thought goes that this addition enables one to do without having an anti-risk condition on knowledge like safety, as the stronger explanatory relation will itself suffice to exclude high levels of epistemic risk.

This proposal doesn't work, however. On the one hand, the cases involving environmental epistemic risk that we have already looked at pose a problem for this view. For not only is the subject's cognitive success significantly attributable to her cognitive agency in these cases, but it is also primarily attributable too, given that nothing intervenes between the skillfully formed belief and the cognitive success. Indeed, in the barn façade scenario, where one is looking at a genuine barn, one is in fact doing everything one needs to do in order to be able to correctly identify it as a barn. Accordingly, robust virtue epistemology is committed to treating the subject as having knowledge, even despite the high levels of epistemic risk involved.²⁰

The problem of environmental epistemic risk illustrates that robust virtue epistemology is too weak, in that it treats some beliefs as being knowledge even though they aren't. The proposal is also too strong, however, in that it excludes some genuine cases from being knowledge. The crux of the matter is that we are often epistemically dependent on others, such as in testimonial cases where our cognitive success is at most only significantly attributable to our cognitive agency. The familiar example in this regard is of someone in good epistemic conditions gaining a true belief via testimony from a knowledgeable informant. Since gullibility is not a route to knowledge, one will expect the manifestation of a certain level of cognitive agency on the part of the subject, such that she is careful about who she asks for information, and suitably critical about the information that she is given. If she gains a true belief as a result, this should suffice for veritic cognitive responsibility, as the cognitive success will be significantly attributable to her cognitive agency. But given that she gained this true belief by, for the most part, trusting a knowledgeable informant, we would surely not regard her cognitive success as being primarily attributable to her cognitive agency, as that would be to negate the important contribution of the informant to this cognitive success.

Robust virtue epistemology is thus both too strong and too weak, and this combination of difficulties makes it hard to respond to these objections, as whatever it does to answer one of the charges will make it more difficult to respond to the other.²¹ In recognizing the problems that face robust virtue epistemology, it becomes clearer why anti-risk virtue epistemology is so attractive.

On the one hand, we've already seen that this proposal can deal with cases of environmental epistemic risk, like the barn façade case. Indeed, it can also deal with related cases like the helpful demon who underwrites the safety of the belief and thereby cancels out the environmental epistemic risk, since in such cases the agent doesn't exhibit safe cognitive responsibility. The crux of the matter is that veritic cognitive responsibility doesn't suffice for knowledge, unlike safe cognitive responsibility.

On the other hand, however, since anti-risk virtue epistemology only demands that that the subject's cognitive agency should play a significant, rather than primary, explanatory role, then it can also deal with scenarios involving epistemic dependence, like the testimonial case that we just looked at. In particular, while these scenarios involve the subject manifesting not merely veritic cognitive responsibility but also safe cognitive responsibility—for if they didn't involve the latter, then they wouldn't be plausible cases of knowledge—they only demand that the subject's cognitive agency plays a significant role, and not also a primary one.

The comparison with robust virtue epistemology also highlights the importance of *how* one strengthens veritic cognitive responsibility to formulate a theory of knowledge. What's required is not to strengthen the explanatory relation *within* veritic cognitive responsibility, but rather to shift from veritic cognitive responsibility to the more demanding safe cognitive responsibility, while keeping the explanatory relation in play fixed. Interestingly, either strengthening preserves the epistemic externalism of the account of knowledge that results. Indeed, even if we opted for a theory of knowledge that opted for both kinds of strengthening—i.e., the shift to safe cognitive responsibility and a strengthening of the explanatory relation—the resulting account of knowledge would still be recognizably externalist since it makes no demand for reflectively accessible reasons. The upshot is that the notion of cognitive responsibility that we are articulating is independent of epistemic internalism. Insofar as part of the motivation for epistemic internalism arises from a desire for our knowledge to demand a robust notion of cognitive responsibility, we are thus in a position to partly accommodate internalist intuitions without thereby embracing epistemic internalism itself.

Nonetheless, there is nothing in anti-risk virtue epistemology that excludes the possibility that most of one's knowledge—indeed, potentially all of it—satisfies an internalist epistemic condition. Indeed, in the most straightforward case, one would expect cognitively developed agents to satisfy this rubric via their possession of reflectively accessible grounds. Anti-risk virtue

epistemology is thus entirely compatible with a moderate epistemic externalism about knowledge.

4. EPISTEMOLOGICAL DISJUNCTIVISM

I just noted that anti-risk virtue epistemology is compatible with agents nonetheless having knowledge that satisfies an internalist epistemic condition. Interestingly, it is in addition compatible with there being certain kinds of knowledge that satisfy such a condition in a particularly full-blooded manner. This is where epistemological disjunctivism comes into the picture. According to this proposal, perceptual knowledge in epistemically paradigm conditions involves reflectively accessible support that is *factive*, in the sense that it actually entails the target proposition that is believed. The epistemically paradigm conditions in question concern one's veridical perception being the result of one's perceptual faculties functioning appropriately in epistemic conditions that are both objectively and subjectively good. This means that not only is the subject in epistemic conditions that are in fact ideal for the subject to gain perceptual knowledge, but that the subject is also appropriately basing her belief on this *factive* reason (e.g., there are no misleading defeaters in play that she is disregarding). Epistemological disjunctivism claims that in these conditions one's perceptual knowledge that *p* can be rationally supported by the *factive* reason that one sees that *p* (where seeing that *p* entails *p*), and where one's rational support is no less reflectively accessible for being *factive*.²²

Such a proposal is controversial, and one can understand why. It has hitherto largely been taken as given in contemporary epistemology that reflectively accessible rational support must be by its nature non-*factive*. This point is often thought to be illustrated by appeal to radical sceptical hypotheses, such as the brain in a vat scenario. The thought goes that one's envatted counterpart, who seems to be experiencing the world in exactly the same way (albeit non-veridically, as she is, unbeknownst to her, trapped in the vat), has exactly the same rational support available to her as one does. If that's right, then of course one's rational support must be non-*factive*, as it is consistent with the target beliefs being false.²³ In contrast, epistemological disjunctivism rejects this line of argument, and maintains that the reflectively accessible rational support available to subjects in 'good' cases (i.e., where conditions are epistemically paradigm) can be different in kind from the reflectively accessible rational support available to subjects in experientially indistinguishable 'bad' cases (e.g., as when one is the victim of a radical sceptical hypothesis). In particular, in the former case the agent can be in possession of *factive* reasons, even though the agent in the latter case only has, at most, non-*factive* reasons available to her. And this is so even though the subject has no discriminative capacity to tell these two cases apart.

Epistemological disjunctivism, if credible, has many attractions, as even its detractors usually acknowledge (their complaint is rather that it isn't credible). Perhaps the chief attraction is the very direct way it offers us of dealing with radical scepticism, at least of a prevalent kind. Just as it is standardly thought that the rational standing for our beliefs is no better than our envatted counterpart, and hence non-factive, so it is widely held that the rational basis for our beliefs, even in the very best epistemic conditions, is compatible with those beliefs being mostly false. Much of the appeal of (at least a certain kind of) radical scepticism follows from this point. If epistemological disjunctivism is true, however, then we can block this rationale for radical scepticism at its source, as in epistemically good conditions the rational standing of our beliefs is clearly very different indeed from our envatted counterparts, and certainly not compatible with the widespread falsity of our beliefs. Epistemological disjunctivism is thus uniquely placed to account for how our reasons can genuinely hook-up with the world, rather than being rational standings that could be completely disconnected to how the world actually is.²⁴

How does the kind of knowledge described by epistemological disjunctivism square with anti-risk virtue epistemology? Let's start with the anti-risk element of the view. If one's basis for belief is a factive reason, then the safety condition is straightforwardly met, as clearly one cannot form a belief on that same basis in close possible worlds and yet believe falsely.²⁵ In particular, there cannot be a high risk of error from forming a belief on this basis, as there is no risk at all. This shouldn't surprise us, given that one needs to be in epistemic conditions that are both objectively and subjectively good in order to enjoy the factive reason in play.

Moreover, notice that the safety of the subject's belief, based on the reflectively accessible factive reason, is also significantly attributable to her manifestation of relevant cognitive agency. This follows from the nature of the epistemic conditions in which the factive reason in question is reflectively available. Recall that the objectively good epistemic conditions that need to obtain involve one's cognitive faculties functioning appropriately in an environment for which they are suited, while the subjectively good epistemic conditions that need to obtain ensure that the subject is appropriately forming her belief on the basis of the reflectively accessible factive reason. It is the former condition which ensures that the subject sees that *p*. But for this to be reflectively available to the subject as a reason for belief, it is important that the latter condition also holds.²⁶ It is the combination of these objective and subjective epistemic conditions that ensures that the safety of the subject's belief is significantly attributable to her manifestation of relevant cognitive agency. The objectively good epistemic conditions entail that one's cognitive agency is a significant factor in enabling one to see that *p* in the first place, and the subjectively good epistemic conditions entail that one is appropriately basing one's belief on this reflectively accessible factive reason (which guarantees the safety of one's belief). It follows that it is not just one's cognitive success that is

significantly attributable to one's cognitive agency, but also one's safe cognitive success.

Indeed, it is arguably the case that when it comes to factive reflectively accessible reasons of the kind that epistemological disjunctivism describes a subject's safe cognitive success will be *primarily* attributable to that subject's cognitive agency. Given the epistemically paradigm nature of the cognitive conditions in which the belief is formed, it is hard to see how other factors could prevent the subject's cognitive agency from playing an overarching explanatory role in her safe cognitive success. In particular, it seems natural to treat the subject's cognitive agency as being primarily responsible for her safe cognitive success, where this is enabled by the optimal cognitive conditions that the subject is operating in. In any case, all that is necessary for our purposes is that epistemological disjunctivism satisfies the weaker explanatory requirement laid down by anti-risk virtue epistemology.

Of course, that epistemological disjunctivism is compatible with anti-risk virtue epistemology doesn't explain why one should find this pairing of views desirable. There is an important rationale in this respect, however, and it concerns what role reflectively accessible reasons can play within a moderate epistemic externalism about knowledge. As we have noted, what is significant about epistemological disjunctivism is that it rejects an orthodox conception of rational support such that one's reasons always fall short of the world, in the sense that they are compatible with widespread falsity in one's beliefs. We can see the relevance of this point to moderate epistemic externalism about knowledge by considering what such an orthodox conception of reasons would contribute to anti-risk virtue epistemology. The satisfaction of safe cognitive responsibility would ensure that one's cognitive agency is appropriately hooked-up with the world. Having, in addition, reflectively accessible non-factive reasons would then be the 'icing on the cake', in that it would enable one to provide a rational story in support of one's beliefs. But notice that such reasons are not themselves anchored in the world at all, due to their non-factive nature, as they are compatible with one's beliefs being massively false. Accordingly, on an orthodox conception of epistemic internalism the satisfaction of an internalist epistemic condition in addition to meeting the anti-risk virtue epistemology rubric thus merely supplies rational support that, as it were, runs parallel to, but is also importantly disconnected from, the safe cognitive responsibility that relates one's beliefs to the world.

In contrast, combining epistemological disjunctivism with anti-risk virtue epistemology leads to a very different kind of proposal. The rational support that one's beliefs enjoy in epistemically paradigm cases of perception is now such that it offers the agent a rational grip on the ways things are in the world, and not merely the kind of cognitively skillful grip on the world represented by safe cognitive responsibility. This is important to preserving the desirability of having reflectively accessible reasons, since if our primary manifestation of cognitive responsibility

is in terms of a skillful responsiveness to the world that is independent of those reasons, then they start to look like mere baubles—nice to have, but serving no deep epistemic purpose. In contrast, by incorporating epistemological disjunctivism into anti-risk virtue epistemology, we can capture a way in which a rational grasp of the world can be manifested from within an externalist account of knowledge.

One concern that one might have about incorporating epistemological disjunctivism within anti-risk virtue epistemology is that factive reasons must impose high cognitive demands on the subject, due to their decisive nature. If that's right, then wouldn't that saddle this putatively 'moderate' epistemic externalism about knowledge with an onerous intellectualism regarding certain core kinds of knowledge?

We noted earlier that questions about the strength of the cognitive responsibility condition—such as whether it demands safe or only veritic cognitive responsibility, or whether the explanatory relation should be cast in stronger or weaker terms—are independent of epistemic internalism, in that on any of these construals one could nonetheless fail to satisfy an internalist epistemic condition. Going in the other direction, it is also true that the strength of the epistemic internalist support that one's belief enjoys can come apart from the extent of cognitive agency on display. In particular, a reflectively accessible factive reason is the strongest kind of internalist epistemic support available, but it doesn't follow that in possessing such a reason one exhibits a superlative level of cognitive agency.²⁷

Indeed, the factive reasons at issue in epistemological disjunctivism are a good case in point in this regard, in that even if they do suffice to ensure that one's safe cognitive success is primarily attributable to one's cognitive agency, this wouldn't be because one has made any great cognitive efforts. In particular, according to epistemological disjunctivism, factive reflectively accessible reasons are not one's reward for exceptional intellectual efforts, but rather occupy a kind of default status in the space of reasons, at least as regards our perceptual engagement with the world.²⁸ It is in virtue of being inculcated into a social fabric of reasons, whereby one cultivates an appropriate sensitivity to relevant abnormal epistemic conditions, that one thereby acquires the practice of presenting factive supporting reasons in epistemically paradigmatic conditions. Epistemological disjunctivism is distinctive in taking this practice at face-value, in opposition to epistemological orthodoxy on this score. In any case, with factive reasons so understood, it is clear why their decisive epistemic status does not require intellectual feats on the part of the subject. The reality is in fact the opposite: factive reasons come easily (at least to subjects suitably embedded in the relevant social practice), it is rather the non-factive reasons that tend to impose higher intellectual demands. That is, it is appropriately responding to departures from epistemically paradigmatic conditions that imposes the higher intellectual burden on agents. Incorporating

epistemological disjunctivism into the heart of anti-risk virtue epistemology thus doesn't pose any danger of leading to an onerous intellectualism about certain core kinds of knowledge.²⁹

5. CONCLUDING REMARKS

Let's try to put together the pieces that we have assembled. The picture that emerges is of a general conception of knowledge—anti-risk virtue epistemology—that demands safe cognitive responsibility. This proposal is a form of epistemic externalism, but it can nonetheless accommodate the internalist insight that knowledge demands cognitive responsibility. It is also compatible with epistemic internalism as being generally applicable to the knowledge possessed by cognitively developed agents. We thus have a moderate epistemic externalism about knowledge.

Then there is epistemological disjunctivism, which we have seen is not only consistent with anti-risk virtue epistemology but when combined with it can capture the sense in which a specifically rational grip on the world can be core to our epistemic standing. Accordingly, by integrating epistemological disjunctivism within anti-risk virtue epistemology we are able to offer a moderate externalist account of knowledge that doesn't merely capture an externalist skill-like conception of cognitive responsibility where reflectively accessible reasons, where possessed, run in parallel to one's knowledge. Instead, we are in a position to explain how the knowledge possessed by cognitively developed agents can fundamentally involve a rational grip on the world that functions in concert with a more general skillful cognitive responsibility.³⁰

REFERENCES

- Baumann, P. (2012). 'No Luck With Knowledge? On a Dogma of Epistemology', *Philosophy and Phenomenological Research* 89, 523-51.
- BonJour, L. (1985). *The Structure of Empirical Knowledge*, Cambridge, MA: Harvard University Press.
- Cassam, Q. (2007). 'Ways of Knowing', *Proceedings of the Aristotelian Society* 107, 339-58.
- Chisholm, R. M. (1977). *Theory of Knowledge* (2nd edn.), Englewood Cliffs, NJ: Prentice-Hall.
- Cohen, S. (1984). 'Justification and Truth', *Philosophical Studies* 46, 279-96.
- Conee, E., & Feldman, R. (2004). *Evidentialism*, Oxford: Oxford University Press.
- Dretske, F. (1969). *Seeing and Knowing*, London: Routledge & Kegan Paul.
- Dutant, J. (2015). 'The Legend of the Justified True Belief Analysis', *Philosophical Issues* 29, 95-145.
- Gettier, E. (1963). 'Is Justified True Belief Knowledge?', *Analysis* 23, 121-23.
- Goldman, A. I. (1976). 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73, 771-91.
- (1979). 'What is Justified Belief?', *Justification and Knowledge*, (ed.) G. Pappas, 1-23, Dordrecht, Holland: D. Reidel.
- (1986). *Epistemology and Cognition*, Cambridge MA: Harvard University Press.
- (1992). 'Epistemic Folkways and Scientific Epistemology', in his *Liaisons: Philosophy Meets the Cognitive and Social Sciences*, 155-78, Cambridge, MA: MIT Press.
- Greco, J. (1999). 'Agent Reliabilism', *Philosophical Perspectives* 13, 273-96.
- (2000). *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*, Cambridge, Cambridge University Press.
- (2009). *Achieving Knowledge*, Cambridge: Cambridge University Press.
- Hetherington, S. (2013). 'There Can be Lucky Knowledge', *Contemporary Debates in Epistemology* (2nd Ed.), (eds.) M. Steup & J. Turri, §7, Oxford: Blackwell.
- Kallestrup, J., & Pritchard, D. H. (2014). 'Virtue Epistemology and Epistemic Twin Earth', *European Journal of Philosophy* 22, 335-57.
- Lehrer, K., & Cohen, S. (1983). 'Justification, Truth, and Coherence', *Synthese* 55, 191-207.
- Le Morvan, P. (2017). 'Knowledge Before Gettier', *British Journal for the History of Philosophy* 25, 1216-38.
- Littlejohn, C. (2009). 'The New Evil Demon Problem', *Internet Encyclopaedia of Philosophy*, (eds.) B. Dowden & J. Fieser, www.iep.utm.edu/evil-new/.
- McDowell, J. (1994). 'Knowledge by Hearsay', *Knowing from Words: Western and Indian Philosophical Analysis of Understanding and Testimony*, (eds.) B. K. Matilal & A. Chakrabarti, 195-224, Dordrecht, Holland: Kluwer.
- (1995). 'Knowledge and the Internal', *Philosophy and Phenomenological Research* 55, 877-93.
- Montmarquet, J. (1993). *Epistemic Virtue and Doxastic Responsibility*, Lanham, MD: Rowman & Littlefield.
- Neta, R., & Pritchard, D. H. (2007). 'McDowell and the New Evil Genius', *Philosophy and Phenomenological Research* 74, 381-96.
- Olsson, E. J. (2012). 'Reliabilism as Explicating Knowledge: A Sketch of an Account', *Epistemology: Contexts, Values, Disagreement*, (eds.) C. Jäger & W. Löffler, 189-202, Berlin: De Gruyter.
- (2015). 'E. J. Gettier and the Method of Explication: A 60 Year Old Solution to a 50 Year Old Problem', *Philosophical Studies* 172, 57-72.
- Plantinga, A. (1993a). *Warrant: The Current Debate*, Oxford University Press, Oxford.
- (1993b). *Warrant and Proper Function*, Oxford University Press, Oxford.
- Pritchard, D. H. (2005a). *Epistemic Luck*, Oxford: Oxford University Press.
- (2005b). 'Scepticism, Epistemic Luck and Epistemic Angst', *Australasian Journal of Philosophy* 83, 185-206.
- (2007). 'Anti-Luck Epistemology', *Synthese* 158, 277-97.
- (2008). 'McDowellian Neo-Mooreanism', *Disjunctivism: Perception, Action, Knowledge*, (eds.) A. Haddock & F. Macpherson, 283-310, Oxford: Oxford University Press.

- (2009a). ‘Apt Performance and Epistemic Value’, *Philosophical Studies* 143, 407-16.
- (2009b). ‘Knowledge, Understanding and Epistemic Value’, *Epistemology (Royal Institute of Philosophy Lectures)*, (ed.) A. O’Hear, 19-43, Cambridge: Cambridge University Press.
- (2011). ‘Epistemological Disjunctivism and the Basis Problem’, *Philosophical Issues* 21, 434-55.
- (2012a). ‘Anti-Luck Virtue Epistemology’, *Journal of Philosophy* 109, 247-79.
- (2012b). *Epistemological Disjunctivism*, Oxford: Oxford University Press.
- (2012c). ‘In Defence of Modest Anti-Luck Epistemology’, *The Sensitivity Principle in Epistemology*, (eds.) T. Black & K. Becker, 173-92, Cambridge: Cambridge University Press.
- (2014). ‘The Modal Account of Luck’, *Metaphilosophy* 45, 594-619.
- (2015a). ‘Anti-Luck Epistemology and the Gettier Problem’, *Philosophical Studies* 172, 93-111.
- (2015b). *Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing*, Princeton, NJ: Princeton University Press.
- (2015c). ‘Risk’, *Metaphilosophy* 46, 436-61.
- (2016a). ‘Epistemic Dependence’, *Philosophical Perspectives* 30, 1-20.
- (2016b). ‘Epistemic Risk’, *Journal of Philosophy* 113, 550-71.
- (2017a). ‘Anti-Risk Epistemology and Negative Epistemic Dependence’, *Synthese* [Online First: <https://doi.org/10.1007/s11229-017-1586-6>].
- (2017b). ‘Knowledge, Luck and Virtue: Resolving the Gettier Problem’, *Explaining Knowledge: New Essays on The Gettier Problem*, (eds.) C. Almeida, P. Klein & R. Borges, 57-73, Oxford: Oxford University Press.
- (2018). ‘Epistemological Disjunctivism and the Biscopic Treatment of Radical Scepticism’, *The Factive Turn in Epistemology*, (ed.) V. Mitova, 15-31, Cambridge: Cambridge University Press.
- (2019a). ‘Epistemological Disjunctivism and Anti-Luck Virtue Epistemology’, *New Issues in Epistemological Disjunctivism*, (eds.) C. Doyle, J. Milburn & D. H. Pritchard, 41-58, London: Routledge.
- (2019b). ‘Epistemological Disjunctivism and Factive Bases for Belief’, *Well-Founded Belief: New Essays on the Epistemic Basing Relation*, (eds.) P. Bondy & J. A. Carter, 235-50, London: Routledge.
- (2020). ‘Anti-Risk Virtue Epistemology’, *Virtue Epistemology*, (ed.) J. Greco & C. Kelp, 203-24, Cambridge: Cambridge University Press.
- Pritchard, D. H., Millar, A., & Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations*, Oxford: Oxford University Press.
- Sainsbury, R. M. (1997). ‘Easy Possibilities’, *Philosophy and Phenomenological Research* 57, 907-19.
- Sosa, E. (1991). *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge: Cambridge University Press.
- (1999). ‘How to Defeat Opposition to Moore’, *Philosophical Perspectives* 13, 141-54.
- (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford: Oxford University Press.
- (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge*, Oxford: Oxford University Press.
- (2011). *Knowing Full Well*, Princeton, NJ: Princeton University Press.
- (2015). *Judgment and Agency*, Oxford: Oxford University Press.
- (2021). *Epistemic Explanations: A Theory of Telic Normativity, and What it Explains*, Oxford: Oxford University Press.
- Williamson, T. (2000). *Knowledge and its Limits*, Oxford: Oxford University Press.
- Wright, C. (2004). ‘Warrant for Nothing (and Foundations for Free)?’, *Proceedings of the Aristotelian Society* (supp. vol.) 78, 167-212.

- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge, Cambridge University Press.
- (1999). 'What is Knowledge?', *Blackwell Guide to Epistemology*, (eds.) J. Greco & E. Sosa, 92-116, Oxford: Blackwell.

NOTES

¹ This is an *accessibilist* construal of the internalist epistemic condition, of a kind defended by Chisholm (1977) and Bonjour (1985, ch. 2), amongst others. This is in contrast to the *mentalist* account of epistemic internalism—offered by, for example, Conee & Feldman (2004)—such that what matters in this regard are the mental states of the subject. Nonetheless, I don't think anything important is lost by focusing on accessibilism in this regard. Mentalists usually grant, after all, that the subject's mental states that are relevant to the epistemic standing of her beliefs are reflectively accessible; their point is just we are interested in such reflective access because we are interested in the mental states being accessed, rather than *vice versa*. Finally, notice that this account of an internalist epistemic condition requires that the subject has reflectively accessible reasons in support of the truth of the target proposition, and not merely in support of her believing it, such as might be supplied by the kind of entitlement strategy advocated by Wright (e.g., 2004) for at least certain kinds of beliefs (in 'cornerstone' propositions). For our purposes, we can set aside how best to locate such proposals within the epistemic externalism/internalism debate.

² At least, this is the standard post-Gettier narrative that gets told (e.g., in contemporary epistemology textbooks). As it happens, I'm inclined to think that the historical situation is much more complex, at least to the extent that insofar as commentators pre-Gettier endorsed anything akin to a theory of knowledge, then it's at least arguable that they were understanding the epistemic condition in an importantly different way to how Gettier (1963) describes it. For some relevant historical discussion of the 'classical' theory of knowledge, see Dutant (2015) and Le Morvan (2017).

³ For the classical defence of process reliabilism, see Goldman (1979; 1986). For a recent defence of the view, see Olsson (2012; 2015).

⁴ In Pritchard (2005*a*; 2005*b*) I call this variety of epistemic luck *reflective epistemic luck*.

⁵ For a key defence of virtue reliabilism (albeit referred to 'agent reliabilism'), see Greco (1999; 2000). For some other proposals that are in the vicinity of virtue reliabilism (albeit not usually described in these terms), see Sosa (1991; 2007; 2009; 2011; 2015; 2021), Goldman (1992), and Plantinga (1993*b*).

⁶ See, especially, Pritchard (2020). See also Pritchard (2016*b*; 2017*a*). As will become clear in a moment, anti-risk virtue epistemology is a variant of *anti-luck virtue epistemology*, which I have defended in earlier work, such as Pritchard, Millar & Haddock (2010, ch. 3) and Pritchard (2012*a*; 2017*b*).

⁷ For some key defences of versions of the safety principle for knowledge, see Sainsbury (1997), Sosa (1999), and Williamson (2000, *passim*). Note that below I apply the general idea behind safety to successful outcomes in general, rather than just to cognitive successes (i.e., true beliefs, for our purposes).

⁸ See Pritchard (2005*a*; 2007; 2012*c*; 2015*a*).

⁹ Indeed, more than this, anti-luck (/anti-risk) epistemology offers us a rationale for understanding the safety principle in a particular way, one that avoids the problems that have been raised against it (e.g., regarding how it accommodates knowledge of necessary truths). It would take us too far afield to get into these issues here, however. For further discussion in this regard, see Pritchard (2005*a*; 2007; 2012*c*; 2015*a*).

¹⁰ For further defence of anti-risk epistemology, and why it is an advance on anti-luck epistemology, see Pritchard (2016*b*; 2017*a*; 2020). For more on the modal accounts of luck and risk that underlie both anti-risk and anti-luck epistemology, see Pritchard (2014; 2015*c*).

¹¹ Inevitably the necessity of safety for knowledge—and the anti-luck/risk intuition that underlies it—has its detractors too. See, for example, Bauman (2012), and Hetherington (2013). See also endnote 20.

¹² This distinction is essentially the same as the distinction between intervening and environmental luck (and their specifically epistemic variants) that I have drawn in earlier work, but with the focus placed instead on the closely related notion of risk. See, especially, Pritchard (2009*a*; 2012*a*) and Pritchard, Millar & Haddock (2010, chs. 2-4).

¹³ For a more precise way of putting this point, see the epistemic twin earth argument discussed in Kallestrup & Pritchard (2014) and Pritchard (2016*a*), which is explicitly designed to capture what is at issue in cases of environmental risk (/luck).

¹⁴ This example is, of course, from Chisholm (1977, 105).

¹⁵ The barn façade case is originally due to Goldman (1976), who credits it in turn to Carl Ginet.

¹⁶ This has been contested, of course—see endnote 20.

¹⁷ This is a variation on the 'Temp' case that I offer in, for example, Pritchard (2012*a*).

¹⁸ For a recent defence of anti-risk virtue epistemology, see Pritchard (2020). See also the defence of anti-luck virtue epistemology offered in Pritchard (2017*b*).

¹⁹ I discuss robust virtue epistemology in a number of places. See especially Pritchard (2009*b*; 2012*a*) and Pritchard, Millar & Haddock (2010, chs. 2-4). For some key defenses of (forms of) robust virtue epistemology, see Sosa (1991; 2007; 2009; 2011; 2015; 2021), Zagzebski (1996; 1999), and Greco (2009).

²⁰ For a further refinement of this objection against robust virtue epistemology, see the epistemic twin earth argument offered in Kallestrup & Pritchard (2014) and Pritchard (2016*a*). Note that some proponents of robust virtue epistemology, such as Sosa (e.g., 2007, ch. 5), bite this bullet and argue that barn façade cases show that there can be lucky knowledge. This is not without obvious costs, as explained in Pritchard (e.g., 2009*a*). Moreover, once we express the point specifically in terms of epistemic risk, then the costs are higher still, as one now needs to contend that knowledge is compatible with high levels of epistemic risk. For more on this point, see Pritchard (2016*b*; 2020).

²¹ For further discussion of this critique of robust virtue epistemology, see Pritchard (2009*b*; 2012*a*; 2016*b*), Pritchard, Millar & Haddock (2010, chs. 2-4), and Kallestrup & Pritchard (2014).

²² Epistemological disjunctivism is rooted in the work of McDowell (e.g., 1995), but the presentation of it here is entirely my own. I develop this proposal in a number of places—see especially Pritchard (2012*b*). See also Neta & Pritchard (2007) and Pritchard (2008).

²³ This is the so-called *new evil demon* intuition. See Lehrer & Cohen (1983) and Cohen (1984). For discussion, see Littlejohn (2009).

²⁴ For a development of this anti-sceptical line, see Pritchard (2015*b*, part 3; 2018; cf. Pritchard 2012*b*, part 3).

²⁵ There are some complex issues raised by the basing relation as it applies to epistemological disjunctivism, though it would take us too far afield to explore them here. For discussion, see Pritchard (2019*b*).

²⁶ This point explains why epistemological disjunctivism rejects the thesis that seeing that *p* entails knowing that *p*, as defended, for example, by Dretske (1969), Williamson (2000), and Cassam (2007). This is because the former can obtain without the latter thereby obtaining (though in such cases it won't be available to the subject as a reflectively accessible reason). See Pritchard (2011; 2012*b*, part 1) for more on this point.

²⁷ At least, it is the strongest kind of internalist epistemic support available insofar as one restricts one's attention to *fallibilist* internalist epistemic support (i.e., epistemic support that is acquired via a fallible cognitive process). The factive epistemic support that is at issue in epistemological disjunctivism is derived via the fallible cognitive process of perception, after all (which is why it is important to keep factivity of epistemic support apart from questions of infallibility). Of course, there might not be any infallible cognitive processes, especially regarding one's epistemic access to facts about an external world, but if there are and they issue internalist epistemic support, then of course they would generate factive rational support of an even more impressive epistemic pedigree.

²⁸ It is an open question whether factive reasons could be extended to other domains. McDowell (e.g., 1994) clearly thinks that they do, though I do not take a stand on this issue here. It will be enough for our purposes that they apply in the perceptual domain.

²⁹ For further discussion of the relationship between epistemological disjunctivism and anti-risk virtue epistemology (at least where the latter is under the guise of anti-*luck* virtue epistemology), see Pritchard (2019*a*).

³⁰ Thanks to Christina Dietz, John Hawthorne, and Luis Oliveira.