THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A sensitive virtue epistemology

Citation for published version:

Collin, J & Bolos, A 2018, 'A sensitive virtue epistemology', *Synthese*, vol. 195, no. 3, pp. 1321-1335.
<https://doi.org/10.1007/s11229-016-1273-z>

Digital Object Identifier (DOI):

[10.1007/s11229-016-1273-z](https://doi.org/10.1007/s11229-016-1273-z)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Synthese

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A Sensitive Virtue Epistemology

We offer an alternative to two influential accounts of virtue epistemology: Robust Virtue Epistemology (RVE) and Anti-Luck Virtue Epistemology (ALVE). We argue that while traditional RVE does offer an explanation of the distinctive value of knowledge, it is unable to effectively deal with cases of epistemic luck; and while ALVE does effectively deal with cases of epistemic luck, it lacks RVE's resources to account for the distinctive value of knowledge. The account we provide, however, is both *robustly* virtue-theoretic and anti-luck, having the respective benefits of both rival accounts without their respective shortcomings. We describe this view here.

Keywords: *virtue epistemology, anti-luck epistemology, sensitivity condition*

1 Overview

Recent work in value epistemology has offered two influential accounts of the value of knowledge: Robust Virtue Epistemology (RVE) and Anti-Luck Virtue Epistemology (ALVE). Central to these accounts is the attempt to discover whether knowledge has, for example, a *distinctive* value over other cognitive states such as mere true belief. One reason for the recent value turn in epistemology is the recognition that merely providing an analysis of knowledge is incomplete. A more robust epistemology would, in theory, tell us not only

what knowledge is, but why obtaining that knowledge is something we ought to value (over and above other cognitive states). Call the attempt to account for the distinctive value of knowledge the *value problem*.

Despite recent interest in the value problem, providing a robust account of the value of knowledge has proved elusive. We argue that there are two primary reasons for this failure. First, attempts to account for the distinctive value of knowledge have failed to adequately deal with cases of knowledge-undermining luck. Second, and in response to the failure noted above, anti-luck accounts which do account for knowledge-undermining luck in return fail to account for the distinctive value of knowledge. Our account, it is argued, avoids these pitfalls and is both robustly virtue-theoretic (and can thus account for the distinctive value of knowledge) and has resources for ruling out knowledge-undermining luck. Before providing the details of the account, it will be important to understand exactly what Robust Virtue Epistemology (RVE) and Anti-Luck Virtue Epistemology (ALVE) claim with regard to the value problem. We begin by considering RVE.

Central to RVE's account of the value of knowledge is the claim that knowledge is a kind of success from ability—it is, in other words, an achievement. The Knowledge-as-Achievement (KA) thesis explains, in an elegant and straightforward manner, that the distinctive value of knowledge is to be accounted for by the distinctive sort of achievement associated with knowledge. It's an achievement where the success comes from the ability of the agent and is distinctive from other kinds of successes where luck could play a contributing role. Put simply, here is the argument:

1. Achievements are finally valuable.
2. Knowledge is a kind of achievement.

Therefore,

3. Knowledge is finally valuable.¹

Three things should be noted about this argument. First, the kind of value associated with knowledge is said to be final value. While we won't discuss the various differences related to value here, final value has non-instrumental value. Though similar to intrinsic value in this regard, final value is distinct from intrinsic value in that it allows for its value to come from some external property. In other words, something is said to be finally valuable because of its relational properties and not merely because of its internal properties. In the case of knowledge, then, its value comes from its relationship to something else valuable—namely, achievements. Second, this argument is also meant to answer a wide variety of value problems, including what Duncan Pritchard calls the *tertiary value problem*—the tertiary value problem being concerned with the distinctive value of knowledge (more on this below). Finally, there is a more general point about this argument that shouldn't be overlooked. Robust virtue epistemologists (RVEs) analyse knowledge in terms of achievements gained through the exercise of cognitive abilities and claim, moreover, that an analysis of this sort is *sufficient* for knowledge, absent an independent anti-luck condition. Only RVEs can make use of the KA thesis to explain the distinctive value of knowledge for the following reasons. If the relevant kind of cognitive achievement is not *sufficient* for knowledge then there will be cases in which an agent can have this sort of achievement without having knowledge. If so, there will be cognitive states which fall short

¹ See Greco (2011, 229-230). As a note, our analysis of knowledge is concerned with only empirical and contingent truths. In the literature, there is a worry that safety and sensitivity conditions cannot account for epistemic luck in the case of necessary truths. We want to avoid this lengthy discussion here and thus limit our claim to empirical and contingent truths. In saying this, though, we do think sensitivity can be naturally modified to deal with necessary truths. See [reference suppressed]

of knowledge but which are also an achievement of this sort; the result being that this sort of achievement cannot account for the *distinctive* value of knowledge.² Making use of the KA thesis to explain the distinctive value of knowledge then requires taking the kind of cognitive achievement at hand to be *sufficient* for knowledge. But this is just what it is to be an RVE. The availability of this explanation of the distinctive value of knowledge is taken to be an advantage of RVE over ‘modest’ virtue epistemologies such as ALVE, which take cognitive achievement to be a necessary but not sufficient condition for knowledge.³

So what exactly is wrong with RVE? To begin, there is a worry concerning the first premise, especially as it relates to the value of easy or wicked achievements. Are wicked achievements, such as those perpetrated by an evil mastermind, really valuable? While this question certainly deserves more attention, it is not one we will consider here.⁴ Instead, our focus is premise 2—the KA thesis. The main problem with the KA thesis, as noted by Pritchard, is that there seem to be cases of cognitive achievements that do not amount to knowledge. If this criticism is successful, then RVE cannot account for the *distinctive* value of knowledge. The details of this argument are spelled out below, but we can note now that Barney-type cases are typical examples where agents seem to meet all the necessary conditions for a genuine cognitive achievement and yet lack knowledge given the element of luck that is involved. The issue for RVE, then, is that it lacks sufficient anti-luck resources to deal with Barney-type cases. The result, Pritchard argues, is that RVE needs to adopt an independent anti-luck condition *and* abandon the claim that *all* knowledge is distinctively valuable. Pritchard of course offers an account that does precisely this in his

² This, then, is the tertiary value problem.

³ That RVE has these unique resources for explaining the distinctive value of knowledge is a point that modest virtue epistemologists have conceded. See (e.g.) Pritchard (2012) and Kallestrup and Pritchard (2012).

⁴ For a discussion on this problem see Pritchard (2009, 5).

"Anti-luck Virtue Epistemology."⁵ We argue here, however, that the criticisms offered by Pritchard do not require an independent anti-luck condition nor do these criticisms require a weakening of the relationship between achievements and knowledge. The advantage of our account is that it maintains the respective benefits of both rival accounts without their respective shortcomings. In other words, it is both *robustly* virtue-theoretic and anti-luck.

In order to see this, let's start by assessing Pritchard's claim that RVE fails to account for the distinctive value of knowledge on the grounds that there are cognitive achievements that do not amount to knowledge.⁶ He illustrates the problem with two examples:

Archie: Archie is a skilled archer who, under normal conditions, successfully hits the target at which he aims. The goal in question is successful because of the ability of Archie. This counts as a genuine achievement. However, on Archie's next shot things are slightly different. Archie selects his target and successfully hits the target because of his skill. Yet, while Archie is successful because of his ability, he is unaware that there are force fields around all the other targets. This means, then, had he selected any other target he would have failed in his goal.

Barney: Barney, unknowingly, is driving through barn façade county. In barn façade county all the barns are made to look like actual barns. Barney, however, happens to see the one barn in barn façade county that is an actual barn. Thus, Barney forms the true belief that he has seen a genuine barn.

⁵ See Pritchard (2012).

⁶ See Pritchard (2010, 50).

The Archie and Barney cases are thought to be relevantly analogous: if one accepts that there is a genuine achievement in the Archie case then one must also accept that there is a genuine (cognitive) achievement in the Barney case. If these cases are indeed analogous, then it will not be the case that knowledge has distinctive value over that which falls short of knowledge. One purported solution, offered by John Greco, is to resist the analogy (2011, 229). Greco argues that the Barney case does not represent success from ability and is therefore not a genuine achievement. Greco conceives of abilities as environment-relative: in barn façade county, Barney doesn't have the ability to distinguish a fake barn from a real barn; as a result, Greco claims, his success in barn façade county is not from ability.

Greco's initial response to these Barney-type cases has generated some serious criticisms⁷ and, as a result, he now offers an updated explanation of environment-relative abilities.⁸ In the updated account Greco still rejects the explicit introduction of modal conditions and maintains that Barney-type cases do not represent genuine achievements. The updated account, however, now appeals to the exercise of intellectual abilities that serve relevant informational needs, where information is needed to enable felicitous action in practical environments. Greco's analysis can be summed as follows:

S knows that p iff:

1. S's believing that p is produced by an intellectual ability of the relevant

⁷ See Pritchard (2010, 34-40).

⁸ Greco (2013).

sort, i.e. of a sort defined by parameters that would serve relevant informational needs;

2. S is in conditions relevant for the exercise of that ability; and

3. S has a true belief because S's belief is produced by an ability of the relevant sort, while in relevant conditions.

Greco's new analysis makes essential use of the notion of actual or potential *practical environments*. In particular the 'relevant sort' of intellectual ability is one that serves informational needs—*viz.* furnishes the agent with information that allows her to achieve her practical aims—in conditions relevant for the exercise of that ability. Abilities, then, are still environment-relative in Greco's new account.

Even with this updated analysis, however, it's still not clear if Greco's account is sufficiently anti-luck (an issue we deal with in more detail below). Another worry is how, on Greco's account, one can, in a non-circular way, pick out the relevant environment. Here is how the problem arises: objective probabilities can only be specified relative to a reference class and any given event belongs to various different reference classes.⁹ Thus, specifying the probability of that event involves first choosing a privileged reference class. If the probability that a car will break down is to be estimated, one must first decide whether to appeal to the frequency of cars breaking down, the frequency of this particular make of car breaking down, the frequency of this particular model of car breaking down, and so on.¹⁰ In

⁹ See Venn (1876) and Reichenbach (1949).

¹⁰ Reichenbach summarizes the role of reference classes in assessing probabilities:

If we are asked to find the probability holding for an individual future event, we must first incorporate the case into a suitable

the present case, the probability that informational needs will be served requires *first* deciding what area around the agent is specially privileged to count as her environment. Brandom develops this point with respect to pure reliabilist theories of knowledge, but it is easy to see how the problem carries over to Greco's analysis:

[S]uppose that Barn Facade County is one of a hundred counties in the state, all the rest of which eschew facades in favor of actual barns. Then ... within the state rather than within the county, our subject's process of perceptual belief formation may be quite reliable ... But then, if the whole country, consisting of fifty larger states, shares the habits of Barn Facade County—so that over the whole country (excepting this one state) facades predominate by a large margin—then considered as a capacity exercised in the country, the very same capacity will count as quite unreliable, and hence as insufficient to underwrite attributions of knowledge. And then again, in the whole world, barns may outnumber facades by a large margin. So considered with respect to that reference class, the capacity would once again count as reliable. And so on. (Brandom (2000, 115-6))

The problem can be pressed in the other direction: if a narrow reference class is considered, such as the environment containing the agent and the barn she happens to be looking at, the capacities she exercises in forming her belief that there is a barn would again count as reliable. In fact, they will count as *maximally* reliable, since the probability the agent will arrive at a true belief through the exercise of her capacities will now be 1. The problem for Greco, then, is that one has to decide what environment the agent is in—what the privileged spatio-temporal region around the agent is—in order to decide whether her

reference class. An individual thing or event may be incorporated in many reference classes, from which different probabilities will result. This ambiguity has been called the *problem of the reference class*. Reichenbach (1949, 374).

informational needs will be served in that environment. Nothing in Greco's account picks out the privileged environment that would allow this process to begin. An important, though not commonly discussed, advantage of safety and sensitivity accounts is that they do precisely this. Safety and sensitivity conditions pick out a sphere of epistemic relevance (one whose size will vary drastically from case to case) and there is no reason at present to think that what constitutes an environment *can* be specified in a non-circular way without appealing to some such modal condition. Environment-relative analyses of knowledge which do not appeal to safety or sensitivity have a *reference class problem*, and in the absence of any reason to think they can solve this problem we ought to be wary any such account.¹¹

Yet, even if environments *could* be pinned down in a way that is both non-circular and makes no implicit appeal to distinct modal conditions, can Greco's new analysis rule out knowledge-destroying epistemic luck? It seems not. The problem is with criterion 1. As Carter, Jarvis, and Rubin (Carter et al.) have argued, what constitutes an agent's 'informational needs' can vary from case to case in such a way that these informational needs can, in the right sorts of conditions, be served even when the agent's belief-forming method is highly unreliable. Importantly, Greco's analysis is *pragmatic* in character: 'informational needs' is understood as picking out that information which is required for successful action. The problem is that there are many cases in which successful action does not require high levels of reliability with respect to belief. One particularly problematic set of cases involves scenarios in which there are high costs associated with not acting, but low costs associated with acting on false information. For creatures who are potential prey,

¹¹ It should be clear that this reference class problem is distinct from the generality problem that pertains to belief forming (see Goldman (1979), Feldman (1995), Conee and Feldman (1998)). That problem becomes apparent when one notices that a belief forming process is a *token* of many different belief forming process *types*, and some principled way of picking out the relevant type is required. A solution to the generality problem will not supply a solution to the reference class problem for environment-relative

belief-forming faculties that generate a high number of false positives of ‘There is a predator in my local environment’ may serve the informational needs of the creatures well, so long as they consistently bring about actions that keep the prey out of harm’s way. Certain medical diagnostic abilities present similar sorts of cases. A disease X is lethal unless treated with drug Y, and there are no harmful side effects to taking Y. Patients who are diagnosed as suffering from X are administered Y. Here, an unreliable method of diagnosis which produced a high number of false positives but zero false negatives would be more conducive to successful action than a highly reliable method of diagnosis which produced some false negatives. The lesson here is that informational needs can be served by beliefs with very low levels of reliability. The problem could be fixed by going hybrid in the manner of Pritchard’s ALVE; appending an anti-luck condition to the virtue-theoretic account. This however would be at the cost of losing the RVE’s ability to explain the unique value of knowledge in terms of achievement.

So, there are serious issues with Greco’s account. There are doubts as to whether it can avoid reference class problems, and, as it stands, it is incapable of ruling out epistemic luck. Further, even if the analysis could be further complicated to deal with these problems—and we see no reason to think it could—there is a more straightforward way of both defending RVE against Pritchard’s objection and ruling out epistemic luck. The approach offered below allows one to endorse the KA thesis, accept the plausible claim that the Barney case involves a cognitive achievement, and circumvent cases of epistemic luck that blight most robust virtue-theoretic accounts of knowledge—all without introducing Pritchard’s independent modal condition.

reliabilist accounts such as Greco’s.

2 Achievement Defined

To see why there is a version of RVE that is not undermined by Pritchard's objection, we need to unpack the notion of achievement. Pritchard argues that the kind of achievement associated with knowledge should be understood in terms of *primary credit*:

(P) Achievements are successes from ability that are primarily creditable to the agent.¹²

(P) is meant to capture the idea that a genuine achievement is one where the success comes about *because of* the relevant abilities of the agent.¹³ Note, there is good reason for spelling out achievements in this way. The notion of achievements as *primarily creditable* to the agent—hereafter, achievements_(pc)—is critical to the KA thesis as it provides a solution to standard Gettier-style cases.¹⁴ The Barney case, however, is not a standard Gettier-style case. The difference in this example is that Barney does in fact observe an actual barn; and, it seems anyway, that he forms the true belief *because of* his cognitive abilities.¹⁵ As Pritchard argues, 'Barney's cognitive success is because of his cognitive ability and so we

¹² Pritchard (2010, 41-42). It's important to note that Pritchard offers this definition in the course of describing the notion of achievement made use of in RVE. Pritchard's ALVE only requires a modest ability condition and rejects the primary credit requirement. RVE, however, is not in a position to reject primary credit given the advantages it has on circumventing Gettier-type problems, such as the Barney case.

¹³ Here we use 'success from ability' and 'success because of ability' synonymously.

¹⁴ Chisholm's sheep case, for example, is a more standard Gettier case. See Chisholm (1977, 105). In this standard case, then, the agent forms the belief that there is a sheep in the field by looking at a dog disguised as a sheep. There is, however, a sheep just out of view. While the belief that 'there is a sheep in the field' is in fact true, the true belief is not *primarily creditable* to the agent. In Pritchard's words, it is not primarily creditable "in the sense that it is to some substantive degree down to her agency that she holds a true belief." Pritchard (2010, 40).

¹⁵ In the end, we reject Pritchard's claim that Barney forms the true belief because of his cognitive abilities (where this notion implies primary credit). While Barney does display some cognitive abilities, he is not primarily creditable in this instance. More on this below.

would, therefore, attribute a cognitive achievement to Barney. That is, his cognitive success in this case is primarily creditable to his cognitive abilities.¹⁶

There are two related questions that need to be answered if we are to respond to Pritchard. First, are all genuine achievements captured by (P)? Put another way, does success from ability entail primary credit? As noted above, there is good reason for RVE to maintain the relationship between a genuine achievement and primary credit. This explains why Greco's response to the Barney case is to deny that it is a representative case of success from ability (i.e., a genuine achievement). If there is no success from ability, there is no genuine cognitive achievement_(pc). Given the noted advantages, then, RVE is wise to insist on the relationship between genuine achievement and primary credit. The problem with this move is that failing to recognize Barney-type cases as genuine achievements is counterintuitive—Barney's gaining a true belief through the exercise of his cognitive abilities does seem to be an achievement *of some sort*—and, as we saw, still isn't sufficient to rule out epistemic luck.

Our suggestion is for RVE is to take Pritchard's analogy seriously and concede that Barney-type cases do represent success from ability. This concession, however, does not necessitate any independent anti-luck condition as Pritchard suggests. The reason for this is that one ought to distinguish *mere* achievement or achievement *simpliciter* from achievement_(pc). Both kinds of achievement are successes from ability, but only the latter is an achievement that is primarily creditable to the agent. We argue that RVE is only committed to the latter. It's not an achievement *simpliciter* which motivates RVE, but achievement_(pc). As such, it is not the case that all achievements are captured by (P):

¹⁶ Pritchard (2010, 51).

achievements_(pc) form a special subset of achievements *simpliciter*, and only achievements_(pc) can account for the distinctive value of knowledge.

Given that the Barney case involves a cognitive achievement of some kind, the second question is whether it involves an achievement_(pc). We will contend that it does not, and hence does not constitute the kind of achievement characterised by (P). An initial problem in addressing this question however is that we do not yet have a clear understanding of what it means for an achievement to be primarily creditable to an agent. Thus, in order to determine whether the Barney case does represent an achievement_(pc), we need to be clear on what (P) entails. Once we have a working definition of primary credit, the ramifications it has for Pritchard's argument become clear—namely, that the Barney case does not represent an achievement_(pc) and the KA thesis is not undermined by Pritchard's objection.

3 The Entanglement Thesis

We are working with a specific notion of achievement in which achievements are successes from ability that are primarily creditable to the agent; achievement_(pc) as opposed to achievement *simpliciter*. The distinction is important because achievement_(pc), we argue, entails sensitivity. This can be seen in the following way. The idea of primary creditworthiness, as we understand it, is supposed to capture the thought, intuitively put, that the achievement is *down to the agent*. In the case of belief this means that getting it right, rather than getting it wrong, is attributable to the activities of the agent herself, and not the

collusion of factors outside of her activities. That is to say, the agent's belief-forming methods are what 'pin down' the correctness of the belief. In those cases in which getting it right is attributable to the agent and her activities, the agent is able to *discriminate* between the truth and falsity of the target proposition.¹⁷ This is a point we will return to shortly, but, in the meantime, if this is not clear, take the (logically equivalent) contrapositive: if an agent is *unable* to discriminate between truth and falsity of the target proposition then her getting it right cannot be primarily attributable to her activities. She is only primarily creditable for gaining a true belief when *her activities* are what pin down the correctness of the belief. An agent then is primarily creditworthy—at least as we understand the phrase here—only if her activities secure her ability to discriminate between truth and falsity in the target case. With respect to belief-formation, this is captured by the following principle:

(*) S deserves to be attributed primary credit for forming a true belief that p, qua *true* belief, only if S's method of belief-formation is able to discriminate between it being the case that p and it not being the case that p.¹⁸

From here, it is only a short step from seeing that primary creditworthiness for getting a belief right has specific modal implications. If one can discriminate between the truth or falsehood of the proposition that p by exercising a particular belief-forming method, then, were one to come to believe that p by that method, one would also fail to believe that p in

¹⁷ Interestingly, Greco (2011, 229) himself recognizes the connection between success from ability and the agent's ability to discriminate between p and ¬p. Greco, though, fails to note the link between discrimination and sensitivity. We discuss this in more detail below.

¹⁸ There is a general worry here about credit and testimony that has been discussed at length in the literature. We won't consider this objection here, but the reader should note Jennifer Lackey's important paper Lackey (2007) on this topic. Lackey argues that in cases of testimony you get knowledge without achievement (testimonial knowledge is too easy for the agent to get credit for the true belief). For a response to Lackey, and one that we think is generally right, see Wayne Riggs (2009).

the closest possible world in which it is not the case that p. This is the sensitivity principle:

Sensitivity: S's belief that p is sensitive if and only if, were that p false, S would not believe that p via the method S actually uses in forming the belief that p.¹⁹

In fact, one of the most attractive features of the sensitivity principle is that it preserves the strong intuition that knowledge of some proposition that p requires the ability to discriminate or distinguish between it being the case that p and it not being the case that p. For Nozick, this was a fundamental reason to endorse sensitivity. For Nozick, and contemporary proponents of sensitivity such as Becker, the discrimination requirement has important implications for how methods must be understood in the sensitivity principle. If we follow Tim Black (2002) in individuating methods externally, by the external causes of belief, then the discrimination requirement is not satisfied. We can see this in the case of BIVs. If external causes are counted within the method, then one's belief that one is not a BIV *is* sensitive, because, were one a BIV, one's external causes of belief would be different. What Nozick and Becker settle on is an internalist characterisation of methods which still respects the externalist nature of the sensitivity principle:

Sensitivity: S's belief that p is sensitive if and only if, were that p false, S would not believe that p via the method_n S actually uses in forming the belief that p.

The subscripted 'n' appended to 'method' indicates that methods are to be individuated

¹⁹ See Dretske (1970) and Nozick (1981) for early presentations of sensitivity principles, and Becker (2007) for a recent

narrowly and in a content-specific way. More specifically, Becker says:

[M]ethods should be individuated by reference to the specific upshot in the agent's experience. That is, which method an agent uses depends on how things *strike* her, and the method should be individuated maximally specifically, in the sense that any feature of how things strike the agent that is causally relevant to belief production ought to be counted as part of the method. (Becker 2012: 91)

An example of this might be: *If you seem to see a hairy, ursine creature with claws, an elongated snout etc. then conclude that there is a bear before you.* Though it's not our goal to defend sensitivity here, two points about the nature of methods_n are briefly worth noting. The first is that Becker is appealing to the content of one's experience in order to parse methods (this is what Becker means above by 'the specific upshot in the agent's experience'). But 'contents' here are understood as non-propositional contents, in order to avoid epistemological internalism. The second point is that this really is an externalist view. Methods are not (or need not be) propositional justifiers since the 'upshot in the agent's experience' need not be understood propositionally. Moreover, the method may not be something the agent himself fully understands and can articulate, or for that matter is even aware of. As such, Becker's sensitivity principle is authentically externalist, whilst retaining the attractive and plausible idea that knowledge requires a discriminative capacity. The link between discrimination and sensitivity (at least, this form of sensitivity) then is not controversial. Not only does discrimination entail Becker's version of sensitivity, this

method-relative formulation of the principle, with a discussion of how to expound the notion of a method.

version of sensitivity was *designed* for precisely that purpose. The ability to discriminate entails that one's methods be sensitive in the Becker way (or something very close to it).

Let's return for a moment to something we briefly discussed above: the link between primary credit and discrimination. First, we are not alone in recognizing this connection; Greco has himself noted that:

[R]elative to Barn Façade County, Barney does not have the ability to perceptually discriminate barns from non-barns. And if he does not have the ability relative to that environment, then his success in that environment is not from ability. (Greco 2011: 229)²⁰

Greco's claim then is that if Barney does not have the ability [to perceptually discriminate barns from non-barns] relative to that environment, then his success is not from ability. And this is the (logically equivalent) contrapositive of our own claim: that if Barney's success is from ability, then Barney does have the ability [to perceptually discriminate barns from non-barns] relative to that environment. We are in agreement with Greco here. If S merits primary credit for arriving at a true belief – if arriving at a true belief is a success from S's abilities – then S must be responsible for her belief being correct; S's belief-forming methods must be what pin down the correctness of the belief. Consider the following example:

S is taking part in a music quiz, where extracts from the compositions of classical

²⁰ See also Greco (2010, 76): 'Henry does not have the ability to tell barns from non-barns relative to the environment he is in. Relative to normal environments, we may assume that Henry can perfectly well discriminate between barns and non-barns.'

composers are played, and S is trying to correctly identify which composer wrote which composition. When S hears polyphonic choral music, S uses the following belief-forming method: *If you seem to hear polyphonic choral music then conclude that you are hearing Rachmaninoff.* The quizmaster is picking between Rachmaninoff, Tallis and Palestrina. S hears an extract from Rachmaninoff's *Liturgy of St John Chrysostom*, forms the belief that the composition is by Rachmaninoff, and answers correctly.

What is at stake is whether S could be primarily creditable for winning a point in the quiz despite not being able to distinguish between the composition being by Rachmaninoff or not. It is clearly wrong in this case to say that S deserves primary credit for getting it right, or that the achievement in this case is down to the activities of S, or that S primarily merits a correct answer in this case, or other cognate claims.

Let's recap. The first main point here is that if S is primarily creditable for correctly believing that p then S is responsible for the correctness of her belief that p. This is simply true in virtue of what we mean by 'primary credit'. The second main point is that if S is unable to discriminate between it being the case that p and it not being the case that p then S is not responsible for the correctness of her belief that p. It is not S's method that pins down the correctness of her belief that p, and one is not responsible for that which is outwith one's auspices. But note that this second point is logically equivalent to 'If S *is* responsible for the correctness of her belief that p, then S *is* able to discriminate between it being the case that p and it not being the case that p'. Here, then, is a clear and simple line of argument from primary credit to discrimination:

Relative to Fake Barn Country, however, Henry does not have that ability.'

1. S is primarily creditable for correctly believing that p. [premise]
2. If S is primarily creditable for correctly believing that p then S is responsible for the correctness of her belief that p. [premise]
3. S is responsible for the correctness of her belief that p. [1, 2, MP]
4. If S is responsible for the correctness of her belief that p, then S is able to discriminate between it being the case that p and it not being the case that p. [premise]
5. S is able to discriminate between it being the case that p and it not being the case that p. [3, 4, MP]

As we have seen the ability to discriminate between it being the case that p and it not being the case that p entails Becker's method-relative form of sensitivity. Hence there is an important but neglected relationship between virtue-theoretic and modal properties, in particular primary creditability entails a particular form of method-relative sensitivity.

This is an *entanglement thesis*: the normative features of knowledge are entangled with the modal features of knowledge; the modal character of knowledge is a result of its normative character. This entanglement thesis is not the important but relatively vanilla observation that virtue-theoretic accounts of knowledge have *some* modal implications, but the recognition that virtue-theoretic accounts of knowledge which make use of (P) entail a very specific and theoretically powerful (if controversial) modal condition on knowledge. Accordingly, in order for an agent to have a true belief that is primarily creditable to her, it needs to be the case that the belief is sensitive. And in Barney-type cases, the belief is clearly

not sensitive. Put simply, Barney does not have the ability to discriminate between p and $\neg p$. Thus, Barney does not have a sensitive true belief that is the result of his cognitive ability and, as a result, the Barney case is not an achievement_(pc). The KA thesis is not undermined by Pritchard's objection.

4 Sensitivity

There is then a point in logical space which allows for an robustly virtue-theoretic account of knowledge which deals with epistemic luck via the sensitivity condition. This account will, of course, differ from both Greco's and Pritchard's since they incorporate (or attempt to incorporate), in their own different ways, a safety condition into their accounts of knowledge.²¹ Is this a problem? In answering this question, it's important to note that any successful account of knowledge needs to describe the sense in which knowledge excludes epistemic luck. With respect to this task, the two major games in town are safety and sensitivity. And while *both* safety and sensitivity are ongoing and interesting research programmes with powerful explanatory resources, each faces a number of objections. Taking just some of the most prominent of these, both safety and sensitivity have been criticised on the grounds that they lead to closure failure—Kvanvig (2004, 2008), Murphy (2005) and Alsepector Kelly (2011) all argue that safety leads to closure failure, whilst Nozick

²¹ A version of the safety principle was first given by Luper (1984), though it is more commonly associated Sosa (1999) who parsed the condition as 'a belief by S that p [is safe] iff: S would believe that p only if it were so that p' (142). More recently, Pritchard (2007, 2008) has endorsed the more sophisticated formulation:

Safety: S's belief is safe if and only if in most nearby possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world, and in all very close nearby possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world, the belief continues to be true.

(1981) held that his sensitivity condition violated the standard closure principle but upheld a more plausible finessed closure principle.²² Bernecker (2012) on the other hand, argues that sensitivity does not violate closure, and Black (2002, 2008) formulates a version of the sensitivity principle in which methods are individuated externally, so that closure is maintained.²³ Comesaña (2005), Baumann (2008), Kelp (2009), Bogardus (2014) and Bogardus and Marxen (2013) all claim that there are instances where knowledge isn't safe; Sosa (1999) and Vogel (2007), on the other hand, claim that there can be insensitive knowledge (though see Becker (2012) for a response to Sosa and Cross (2007) for a response to Vogel). McEvoy (2009) and Dodd (2012) have also argued that safety cannot handle lottery cases.

As one can see, then, the issues surrounding safety and sensitivity are ongoing. And defending sensitivity principles against all the objections leveled against them is beyond the scope of this paper (or, for that matter, any single paper).²⁴ Like any active research programme, and like its cousin safety, it will be subject to a number of criticisms and a number of attempts to circumvent these criticisms, or otherwise show that they are unfounded. Our approach to the sensitivity condition itself is modest: whether it stands *at the end of inquiry* is not something to be settled here. What we *do* claim however is that our proposal has a unique advantage over its competitors in that it is both *robustly* virtue-theoretic and anti-luck; and, moreover, that sensitivity should *at the very least* be kept on the table, as a going concern, and is not obviously subject to knock-down objections. It

²² See Baumann (2012) for discussion and for an argument that the standard closure principle licenses bootstrap-ping and should be rejected.

²³ Perhaps for sociological, rather than philosophical reasons, closure failure is widely perceived to be a problem for sensitivity alone. This perception is not reflected in the work carried out in this area.

²⁴ Though we do not deal with objections to sensitivity here, we do offer a positive defense. We note, for example, the particular advantages sensitivity has over safety with respect to the value problem. This is consistent with a central aim of the paper—to note the connection between sensitivity and virtue epistemology. In the end, the claim regarding sensitivity is merely conditional.

is, in other words, a live research option. As a result, the sensitive virtue epistemology we have described here is a legitimate rival to standard RVE and ALVE accounts.

5 The Benefits of a Sensitive Virtue Epistemology

There is good reason to explicate knowledge in terms of primary creditworthiness, since doing so, as was noted above, both addresses standard Gettier problems and provides an attractive response to the value problem. Doing so also, we claim, vindicates our contention from section 1. One can hold that the Barney case is a cognitive achievement, but because it is not the sort of cognitive achievement relevant to knowledge it is not impacted by Pritchard's worry. Furthermore, primary creditworthiness *entails* sensitivity, so robust virtue epistemologists who explicate their account of knowledge in terms of primary creditworthiness *must* endorse sensitivity if they are to avoid incoherence. The result of this is that this robust virtue-theoretic account has anti-luck credentials built-in. Neither Barney cases nor, yet trickier, epistemic twin earth cases²⁵ will cause special problems for this account that are not faced by analyses of knowledge that include independent modal anti-luck conditions. This gives the preceding account distinct advantages over other robust virtue theoretic accounts and hybrid virtue-theoretic, anti-luck accounts. The former can endorse the KA thesis, but are subject to counterexamples involving epistemic luck, whilst the latter avoid these counterexamples at the cost of abandoning the KA thesis. No such choice need be made here.

Two other claims follow from the entanglement thesis. The first is irenic: Pritchard

²⁵ Cf. Kallestrup and Pritchard (2012).

and the robust virtue epistemologist are not so far apart as one may have supposed. Pritchard *conjoins* a virtue-theoretic condition with an ‘independent’ modal anti-luck condition to forge his account of knowledge. We claim, on the other hand, that at least one virtue-theoretic condition *entails* a modal anti-luck condition. This is to the good, since unless robust virtue epistemologists simply *deny* that knowledge is incompatible with veritic luck, they *must* endorse an anti-luck condition, whether it is made explicit or is implicit in their preferred virtue-theoretic conditions.²⁶ The second, related, claim is that, whilst the virtue theoretic account of knowledge is not subject to the criticism discussed here, it is also a larger target than virtue-theoretic epistemologists might have realised: given its entailment of sensitivity, objections to sensitivity are, by that fact, objections to the virtue-theoretic account. Some epistemologists will see an entailment of sensitivity as a perquisite, others will not; but as we noted earlier, sensitivity, like safety, is research programme worth pursuing. As a result, there is a place in logical space for a RVE epistemology with the benefits, but not the drawbacks, of Greco’s account, and the benefits, but not the drawbacks, of Pritchard’s hybrid ALVE. A final feature of this account deserves note. As Pritchard (2012) has emphasized, two intuitions govern our thinking about knowledge. One is an anti-luck intuition: in cases of knowledge, the truth of one’s belief cannot just be a matter of luck. This is why Gettier, and cognate, cases do not constitute knowledge.

The second intuition is that in cases of knowledge, the cognitive success at hand is the product of cognitive abilities. There Pritchard contends that these two intuitions are disjoint. Our account draws the opposite conclusion; that there is a deep theoretical unity between the normative and the modal aspects of knowledge. That a sensitive virtue

²⁶ Irenic, but not utterly ecumenical, since Pritchard rejects the sensitivity condition in favor of safety (Pritchard (2005)).

epistemology has this unifying explanatory power is a uniquely attractive feature.

References

- Alspector Kelly, M. (2011, June). Why safety doesn't save closure. *Synthese* 183(2), 127–142.
- Baumann, P. (2008). Is knowledge safe? *American Philosophical Quarterly* 45(1), 19–30.
- _____. (2012). Nozick's defense of closure. In K. Becker and T. Black (Eds.), *The Sensitivity Principle in Epistemology*. Cambridge University Press.
- Becker, K. (2007). *Epistemology Modalized*. Routledge Studies in Contemporary Philosophy. Routledge.
- _____. (2012a). Methods and how to individuate them. In K. Becker and T. Black (Eds.), *The Sensitivity Principle in Epistemology*. Cambridge University Press.
- _____. (2012b). Sensitivity, safety and closure. *Acta Analytica* 27, 376–381.
- Black, T. (2002). A moorean response to brain-in-a-vat skepticism. *Australasian Journal of Philosophy* 80(2), 148–63.
- _____. (2008). Defending a sensitive neo-moorean invariantism. In V. Hendricks and D. Pritchard (Eds.), *New Waves in Epistemology*. Basingstoke: Palgrave Macmillan.
- Bogardus, T. (2014). Knowledge under threat. *Philosophy and Phenomenological Research* 88(2), 289–313.
- Bogardus, T. and C. Marxen (2013). Yes, safety is in danger. *Philosophia* 2, 1–14.
- Brandom, R. (2000). *Articulating Reasons: An Introduction to Inferentialism*. Harvard University Press.

- Carter, J. A., B. W. Jarvis, and K. Rubin (2015). Varieties of cognitive achievement. *Philosophical Studies*.
- Chisholm, R. M. (1977). *Theory of Knowledge* (2nd ed.). *Foundations of Philosophy*. Prentice Hall.
- Comesaña, J. (2005). Unsafe Knowledge. *Synthese* 146(3), 395–404.
- Conee, E. and R. Feldman (1998). The generality problem for reliabilism. *Philosophical Studies* 48, 1–29.
- Cross, T. (2007). Comments on Vogel. *Philosophical Studies* 134, 89–98.
- Dodd, D. (2012). Safety, skepticism, and lotteries. *Erkenntnis* 77(1), 95–120.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy* 67(24), 1007–23.
- Feldman, R. (1995). Reliability and justification. *The Monist* 68, 159–174.
- Goldman, A. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and Knowledge*, pp. 1–23. Dordrecht: D. Reidel.
- Greco, J. (2011). The value problem. In *The Routledge Companion to Epistemology*. Routledge.
- _____. (2013). A (different) virtue epistemology. *Philosophy and Phenomenological Research* 85(1), 1–26.
- Kallestrup, J. and D. Pritchard (2012). Virtue epistemology and epistemic twin earth. *European Journal of Philosophy* 21(2).
- Kelp, C. (2009). Knowledge and safety. *Journal of Philosophical Research* 34, 21–31.
- Kvanvig, J. (2004). Nozickian epistemology and the value of knowledge. *Philosophical Issues*

14(1), 2001–18.

_____. (2008). Critical notice of pritchard's epistemic luck. *Philosophy and Phenomenological Research* 77, 272–284.

Lackey, J. (2007). Why we don't deserve credit for everything we know. *Synthese* 158(3), 345–61.

Luper, S. (1984). The epistemic predicament: Knowledge, nozickian tracking, and scepticism. *Australasian Journal of Philosophy* 62(1), 26–49.

McEvoy, M. (2009). The lottery puzzle and pritchard's safety analysis of knowledge. *Journal of Philosophical Research* 34, 7–20.

Murphy, P. (2005). Closure failures for safety. *Philosophia* 33(331-34).

Nozick, R. (1981). *Philosophical Explanations*. Harvard University Press.

Pritchard, D. (2005). *Epistemic Luck*. OUP.

_____. (2007). Anti-luck epistemology. *Synthese* 158(3), 277–297.

_____. (2008). Knowledge, luck and lotteries. In V. Hendricks (Ed.), *New Waves in Epistemology*. Palgrave Macmillan.

_____. (2009). Knowledge, understanding and epistemic value. In *Epistemology*. Cambridge University Press.

_____. (2010). Knowledge and final value. In *The Nature and Value of Knowledge: Three Investigations*. Oxford: OUP.

_____. (2012). Anti-luck virtue epistemology. *Journal of Philosophy* 109(3), 247–279.

Reichenbach, H. (1949). *The Theory of Probability*. University of California Press.

Riggs, W. (2009). Two problems of easy credit. *Synthese* 169, 201–216.

Sosa, E. (1999). How to Defeat Opposition to Moore. *Noûs* 33(13), 141–153.

Venn, J. (1876). *The Logic of Chance* (2 ed.). Macmillan and co.

Vogel, J. (2007). Subjunctivitis. *Philosophical Studies* 134(1), 73–88.