



Carter, J. A. and Pritchard, D. (2017) Cognitive bias, scepticism and understanding. In: Grimm, S. R., Baumberger, C. and Ammon, S. (eds.) *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*. Routledge: New York, pp. 272-292. ISBN 9781138921931

This is the author's final accepted version.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/136260/>

Deposited on: 06 February 2017

Enlighten – Research publications by members of the University of Glasgow
<http://eprints.gla.ac.uk>

For *Explaining Understanding: New Perspectives from Epistemology and the Philosophy of Science*,
(eds.) S. Ammon, C. Baumberger, C. Beisbart, G. Brun & S. Grimm, (Routledge, *forthcoming*).

COGNITIVE BIAS, SCEPTICISM AND UNDERSTANDING

J. Adam Carter & Duncan Pritchard
University of Edinburgh

ABSTRACT. In recent work, Mark Alfano (2012; 2014) and Jennifer Saul (2013) have put forward a similar kind of provocative sceptical challenge. Both appeal to recent literature in empirical psychology to show that our judgments across a wide range of cases are riddled with unreliable cognitive heuristics and biases. Likewise, they both conclude that we know a lot less than we have hitherto supposed, at least on standard conceptions of what knowledge involves. It is argued that even if one grants the empirical claims that Saul and Alfano make, the sceptical conclusion that they canvass might not be as dramatic as it first appears. It is further argued, however, that one can reinstate a more dramatic sceptical conclusion by targeting their argument not at knowledge but rather at the distinct (and distinctively valuable) epistemic standing of *understanding*.

1. SCEPTICAL ARGUMENTS: SOME DISTINCTIONS

Sceptical arguments in epistemology usually take a distinctive form: they raise doubt that some epistemic standing—typically, knowledge, though not always—is as prevalent as it is ordinarily taken to be. To the extent that such arguments succeed, the upshot is that we know, justifiably believe (etc.), much less than we supposed, and thus our total epistemic position is left worse off than anticipated. It will be useful to note three ways in which we might individuate varieties of scepticism.

First, scepticism can be *radical* or *local* depending on the scope of its ambitions. At one extreme, one can imagine a form of scepticism that maintained that none of our beliefs could ever have a positive epistemic standing of any kind. This would be radical scepticism indeed. But equally one can also imagine sceptical arguments that merely target a particular kind of epistemic standing, such as our knowledge of the past, and hence which are local.

Second, sceptical arguments can be individuated in terms of which specific epistemic standing is being targeted. Ordinarily, we think of scepticism as targeting knowledge—either radically (e.g., Cartesian scepticism) or more locally (e.g., scepticism about the past)—but scepticism can also potentially target epistemic states other than knowledge, such as justification.

Finally, third, scepticism can be motivated either empirically or on purely *a priori* grounds. In the contemporary literature at least, sceptical arguments tend to be of the latter variety. Consider, for example, a standard form of radical scepticism about our knowledge of the external world that appeals to radical sceptical hypotheses. Very roughly (the details need not concern us), this proceeds by arguing that since we cannot exclude these sceptical scenarios, so we lack much, if not all, of the knowledge of the external world that we take ourselves to possess. There is no empirical commitment being made here—the sceptic is not, for example, claiming that we are victims of sceptical scenarios, or even that it's likely that this is the case. They are merely registering this possibility, and the fact that this is something that we need to exclude if we are to have knowledge of an external world, and then noting that we cannot exclude it (and that from this sceptical conclusions follow).

In contrast, an empirically-orientated scepticism will appeal to specific empirical claims to motivate the target sceptical conclusion. Inevitably, scepticism of this variety will tend to be local rather than radical, since otherwise the scepticism would undermine the empirical basis cited in its favour, and so would be self-defeating. But scepticism of this sort can be no less intellectually devastating if successful, depending on the epistemic standing in question. If, for example, it were shown that we have far less knowledge of X than we hitherto supposed, and yet knowledge of X is something that we highly value, then this would be a troubling form of scepticism even despite being local rather than radical.¹

In what follows, a novel variety of scepticism will be motivated. It will be empirically-orientated, in that it appeals to some of the latest empirical psychology on cognitive biases. As such, it will be a local form of scepticism. But, as we will see, the sceptical conclusion being canvassed is no less intellectually troubling in virtue of the scepticism being local.²

2. BIAS-DRIVEN SCEPTICISM: TWO STRANDS

In recent work, Jennifer Saul (2013*a*) and Mark Alfano (2012; 2014) have motivated different versions of a very provocative strand of scepticism about at least some kinds of human knowledge. Both Saul

and Alfano appeal to empirical premises, supported by the latest psychology, to the effect that, as Saul (2013*a*, 243) puts it, we ‘have very good reason to suppose that we are systematically making errors caused by our unconscious biases’. Furthermore, both authors take the fallout of this observation about bias and error to threaten the scope of knowledge that we can lay claim to. While in very close alignment in terms of their central objective, however, they differ with respect to which kinds of biases they primarily focus on. They also differ in terms of how they reason from the empirical facts they cite to local sceptical conclusions.

2.1. SAUL ON IMPLICIT BIAS AND KNOWLEDGE

A bias, in the most general sense, is a disposition, implicit or explicit, to reach a particular kind of conclusion or outcome—in the kind of case we’re interested in, the outcomes will be representational. Call these *cognitive* biases.³ One kind of cognitive bias consists in ‘unconscious tendencies to automatically associate concepts with one another’, tendencies which, as Saul (2013*a*, 144) notes, can and often do lead to some ‘disturbing errors’. The most studied and perhaps most dangerous such associations involve, as Jules Holroyd and Joseph Sweetman (2015, 4) put it, ‘relating cognitive (including affective) content, that is mental representations or associations, to behavioural dispositions, in a particular context’. Examples include associating implicitly certain context-specific performance behaviours (e.g., academic performance, athletic performance, intelligence) with concepts such as racial category, religion, and so on. Call this class of cognitive bias, which we are typically oblivious to, *implicit bias*.⁴

Saul gives a number of examples which help to show just how prevalent implicit biases are in everyday judgments, particularly regarding our judgments about individuals from stigmatised groups. Consider, for example, our perception of CV quality. Saul notes that in cases where the experimenter holds fixed all items on a CV, switching out only the names at the top, what is found is that:

[T]he same CV is considered much better when it has a typically white rather than typically black name, a typically Swedish rather than a typically Arab name, a typically male rather than typically female name, and so on.⁵ (Saul 2013*a*, 244)

As Saul points out, such judgments are influenced by factors to do with social category that should be entirely irrelevant. Another striking example of implicit bias Saul draws attention to is ‘shooter bias’, which skews one’s perceptual awareness. Shooter-bias cases reveal that a given ambiguous object is significantly more likely to be perceived as a gun rather than as something ‘innocent’ (e.g., a phone) when held by a young black or Muslim man rather than when held like by a young white man.⁶

Saul's route from observations about implicit bias to a form of scepticism about knowledge proceeds by first noting that facts about implicit biases show that 'we have very good reason to suppose that we are systematically making errors caused by our unconscious biases related to social categories.' (Saul 2013a, 250) If such judgments related to social categories were rare, then this would motivate only a very localised form of scepticism about knowledge. According to Saul, however, such judgements are widespread:

The problem starts to become vivid when we ask ourselves when we should be worried about implicit bias influencing our judgments. The answer is that we should be worried about it whenever we consider a claim, an argument, a suggestion, a question, etc from a person whose apparent social group we're in a position to recognize. Whenever that's the case, there will be room for our unconscious biases to perniciously affect us. (Saul 2013a, 255)

One pervasive such arena is that of testimony, where we are subject to unconscious biases about testifiers' credibility. And what goes for our assessments of testimony applies *mutatis mutandis* to other social-epistemic activities, such as considering questions, judging arguments, listening to contributions, and so on. It thus follows that Saul has potentially identified an empirical basis for scepticism about a reasonably wide class of knowledge.⁷

Interestingly, Saul further argues that the extent to which bias-related doubt suggests we cannot properly trust, as she puts it, our 'knowledge-seeking faculties' generates a more powerful consideration in favour of scepticism than does the traditional *a priori* approach that appeals to radical sceptical hypotheses. She writes:

[A] standard response [to traditional radical sceptical scenarios] is that these worries should not grip us, because we have no reason at all to suppose that these possibilities obtain. Doubt induced by implicit bias *is unlike this*: we have very good reason to suppose that we are systematically making errors caused by our unconscious biases related to social categories. (Saul 2013a, 241)

Saul's reasoning here is somewhat odd, however, in that sceptical arguments that trade on these hypotheses aren't claiming that there is any reason for thinking that these scenarios obtain, or are even likely to obtain. As we saw above, the point is rather that our inability to exclude them is meant to have radical sceptical consequences, regardless of whether such scenarios are in fact likely or actual. If that point is sound, then it holds regardless of whether we have a rational basis for supposing that they do obtain. In that case, there wouldn't be any intellectual comfort in noting that one has no rational basis for thinking that they do obtain, since one lacks knowledge regardless.

Still, the wider point that Saul wants to make here, about how the form of scepticism she is proposing is very different from traditional forms of scepticism in virtue of being empirically

motivated, is certainly sound. Returning to our three-fold classification of sceptical arguments above, what Saul is offering us is a form of scepticism that is:

- (i) local, albeit such that it targets a wide enough class of knowledge to be intellectually disturbing.
- (ii) directed at knowledge;
- (iii) empirically motivated; and

Moreover, we take it that the point Saul is getting at by comparing the error-scenarios involving cognitive bias to radical sceptical hypotheses is that the sceptical claim in play here is not merely that more of our beliefs are false than we were hitherto aware of (and hence that we know less than we thought for this reason). Rather, there is a much more disturbing claim in play, which is that *even if* our beliefs are not formed as a result of bias—and, indeed, even if those beliefs are true—the widespread and unconscious nature of such bias means that many of our beliefs do not amount to knowledge. Basically, Saul’s thought is that given that we are ordinarily willing to concede that our inability to rationally exclude radical sceptical scenarios can potentially undermine knowledge (scenarios which we have no reason to think obtain), so we should be very disturbed by the *actual presence* of cognitive bias in our judgements. Unless we are able to rationally exclude the possibility that our judgements within a relevant domain are subject to this bias, then they do not amount to knowledge *even if they are true and unaffected by bias*.

2.2. ALFANO ON EPISTEMIC SITUATIONISM AND KNOWLEDGE

In contrast to Saul, Alfano’s (2012; 2014) brand of scepticism finds its source of motivation in an influential strand of thinking in the literature on moral psychology known as *situationism*. According to situationism, there is a body of empirical literature which demonstrates that ‘extra-agential’ factors—such as ambient light levels and sounds, ambient smells and mood depressors, the presence of bystanders, hunger levels and so on—can influence our actions and judgments to a surprising extent. John Doris (1998; 2002) and Gilbert Harman (1999; 2000; 2003) have reasoned from these kinds of empirical observations to the conclusion that virtue ethics is empirically inadequate. In particular, the idea is that the virtue ethicist’s postulation of stable character traits—i.e., moral virtues and vices—in the service of explaining moral (and immoral) behaviour is empirically undermined by the extent to which extra-agential factors, including ones of which the subject may be consciously unaware, seem to be doing much of the relevant explanatory work. We thus get a kind of scepticism, albeit not of one the epistemological kind (i.e., which specifically targets epistemic standings) that interests us here.

This is where Alfano comes in, as he has argued that the general situationist line taken by Doris and Harman against virtue ethics can be extended to apply to epistemology. In particular, his target is a view widely held in epistemology known as *virtue epistemology*. This holds that knowledge is acquired via the manifestation of *epistemic virtues*, where this is a class of belief-forming processes which includes cognitive abilities, cognitive faculties, and intellectual virtues. Accordingly, if situationism can undermine virtue ethics by showing that our ethical virtues do not play the explanatory role in our ethical behaviour that the view proposes, then in principle the very same situationist critique can be applied *mutatis mutandis* to virtue epistemology to show that our epistemic virtues do not play the explanatory role in our epistemic behaviour (e.g., in the way we form beliefs) that the view proposes. The result would be a dilemma: either virtue epistemology must be rejected, or else one must embrace the sceptical consequence that we know a lot less than we hitherto supposed.

Alfano's (2012) redeployment of the situationist argument against the *virtue responsibilist* brand of virtue epistemology is relatively straightforward. Virtue responsibilist approaches to knowledge (e.g., Zagzebski 1996) insist that knowledge is the product of intellectual character virtues (e.g., intellectual honesty, open-mindedness, etc.), where these virtues involve distinctive motivational traits (e.g., a desire for truth) and are reasons-responsive. Given the tight commonalities between intellectual character virtues and moral character virtues, the Doris-Harman line would seem to apply in equal measure to virtue responsibilists, who (like virtue ethicists) ultimately need stable intellectual character traits so described to play important explanatory roles.

But many virtue epistemologists do not endorse the responsibilist model, but rather the *virtue reliabilist* model (e.g., Sosa 2007; 2009; Greco 2010). On this view, the traits doing the explanatory work in cases of knowledge acquisition are not exclusively conceived of as robust character traits like intellectual virtues, but also include cognitive faculties and cognitive abilities, such as the cognitive traits involved in the basic perception of one's immediate environment. Such faculties and abilities can, on the virtue reliabilist model, be very specific and needn't involve any distinctive motivational states or responsiveness to rational considerations. All that is required is that such abilities are in fact reliable (in the relevant conditions) and that they are suitably integrated with the subject's other character traits so as to count as part of the overarching cognitive character of the subject.⁸

Obviously, the sense in which situational factors would undermine a virtue reliabilist approach to knowledge is not as clear-cut as it is on the responsibilist model. Alfano (2014) himself is careful to limit his challenge (with respect to virtue reliabilism) to *inferential* knowledge. He concedes that *non-inferential* knowledge will, on the virtue reliabilist model, be immune to the brunt of the situationist

critique. As his argument runs, however, extra-agential factors play a significant *enough* role in explaining our cognitive successes that it is implausible to regard cognitive faculties and abilities as playing the substantial explanatory role virtue reliabilists claim they play in cases of inferential knowledge. So even here, insofar as one retains one's commitment to virtue epistemology, then a local form of scepticism about knowledge is the result.

In support of his attack on virtue epistemology, Alfano (2014, 110) draws attention to the empirical literature on cognitive biases, such as the *availability* and *representativeness* heuristics (Tversky & Kahneman 1973; 1974). The availability heuristic leads people to expect that the probability of (say) an event is positively correlated with the ease by which the event in question can be brought to memory. A simple study that illustrates this idea—due to Amos Tversky and Daniel Kahneman (1973)—goes as follows. Participants were asked how likely it is that a random word taken from an English text starts with a 'K', or that 'K' is the third letter. As it happens, it is approximately twice as likely that a given word has 'K' as the third rather than first letter, but since participants tend to find it easier to think of words that start with 'K' than words that have 'K' as the third letter, they tend to judge that it is much more likely that the random word chosen will start with a 'K'.

One area where the availability heuristic has been especially well-documented is in the literature on risk perception. The more exposure one has to some perceived risk (e.g., Ebola) the higher one is likely to estimate the probability of the risk. For example, one who goes to a restaurant where conversations about Ebola can be overheard is more likely, *ceteris paribus*, to estimate the risk-probability of Ebola than one who goes to a different restaurant where no such conversations happen to be going on.⁹

One factor that is positively correlated with ease of recall (which then influences probability estimates) is how *recently* one has been exposed to some idea or concept. In one study (deTurck *et al* 1990), mock jurors were significantly more inclined to regard a testifier as deceptive if the testifier lied *after* telling the truth, than were the testifier to lie and then tell the truth. In general, it seems that information encoded from our most *recent* encounters typically are (unconsciously) afforded disproportionate weight.

Another well-studied cognitive bias Alfano draws attention to is the *representativeness bias* (Tversky & Kahneman 1974). Consider the famous case of Linda. 'Linda is thirty-one years old, single, outspoken and very bright. She majored in philosophy. As a student she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.' Participants, in light of this information, were asked to assess the probability that various attributes are

true of Linda. Among them were that (i) Linda is a bank teller; and (ii) Linda is a bank teller and is active in the feminist movement. 89% of participants rated (ii) more probable than (i), even though this is probabilistically impossible.¹⁰ What explains this mistake, according to Tversky and Kahneman (1974), is our instinctive use of heuristics, such as ‘stereotyping’, or judging likelihood/frequency of something on the basis of its perceived resemblance to the stereotype of the item or individual in question.¹¹

Putting this all together, cognitive biases (such as availability and representativeness) are just some of the many biases that we unconsciously employ all the time, and which illustrate just how highly sensitive our judgments are to things that seem entirely external to the kinds of cognitive abilities we regard as germane to knowledge acquisition. So insofar as knowledge is the product of intellectual virtue, or is gained inferentially via the use of cognitive virtues and faculties, then we have much less knowledge than we suppose.¹² That’s a pretty broad class of knowledge, at least if virtue epistemology is correct (i.e., and that knowledge is not generally acquired in ways that have nothing to do with one’s exercise of epistemic virtue).¹³

Alfano has thus motivated a form of scepticism that is, with reference to our three-fold taxonomy in §1, of essentially the same kind as that advanced by Saul. In particular, it is:

- (i) local, albeit such that it targets a wide enough class of knowledge to be intellectually disturbing.
- (ii) directed at knowledge;
- (iii) empirically motivated; and

Furthermore, notice that Alfano’s situationism-driven scepticism is also akin to Saul’s bias-driven scepticism in another key respect, although he is less explicit about this. We noted above that Saul’s overarching point seemed to be not merely that some of our beliefs that we thought were true turned out to be false, and so not in the market for knowledge, on account of being the result of bias. Rather her sceptical worry was the more disturbing one that once we grant that we are regularly subject to unconscious bias, then insofar as we cannot rationally exclude this possibility it follows that our beliefs do not amount to knowledge even when true and not infected by bias.

Although it is less clear in this case, Alfano at least seems to be making a parallel kind of point—*viz.*, that once we grant that our belief-forming practices are unconsciously affected by situational factors in the ways he has drawn attention to, then our general practice of attributing our cognitive successes to our own exercise of epistemic virtue is called into question. In particular, *even if* one’s belief is the product of epistemic virtue, so long as one is unable to rationally exclude the possibility that situational factors are playing an explanatory role in one’s belief-formation within a

relevant domain, then that belief cannot amount to knowledge. And that is a far more disturbing sceptical claim. In any case, henceforth we will take Alfano to be endorsing, in line with Saul, the stronger sceptical conclusion.¹⁴

3. RE-EVALUATING THE BIAS-DRIVEN SCEPTICAL CHALLENGE

In what follows we will grant the empirical basis of the scepticism offered by Alfano and Saul. In particular, we are granting that the empirical support offered by Alfano and Saul demonstrates at least this much: that often our beliefs in the relevant domains are infected by cognitive bias—sometimes to the extent that they are primarily due to cognitive bias—with the result that some of these beliefs, so formed, are false (call this the *bias thesis*). In granting the bias thesis, we are thereby also granting that Saul and Alfano have at least motivated a limited kind of scepticism. After all, false beliefs are not even in the market for knowledge, and most epistemologists (and certainly virtue epistemologists at any rate) would surely agree that even true beliefs that are primarily the result of cognitive bias are not in the market for knowledge either. It thus follows from the bias thesis that, at least in certain domains, we know a lot less than we hitherto supposed. Call this *weak bias-driven scepticism*.

Alfano and Saul clearly want to motivate something stronger than weak bias-driven scepticism, however, and it is this more demanding form of scepticism that is our interest from here on in. We think we can delineate two kinds of scepticism in this regard. According to the first—which we will refer to as *intermediate bias-driven scepticism*—the thought is that even where epistemic virtues are genuinely in play in the formation of a true belief, so long as cognitive bias is also playing a role in this regard, even a relatively minor one, then this will infect the epistemic standing of the belief and thereby prevent it from counting as knowledge. An even wider class of knowledge is thus called into question compared with weak bias-driven scepticism. According to the second type of scepticism—which we will refer to as *strong bias-driven scepticism*—even when true beliefs in the relevant domain are not infected by cognitive bias, since we are unable to rationally exclude the possibility that cognitive bias is present, it nonetheless follows that we lack knowledge. A wider still class of knowledge is thus undermined.

How persuaded should we be of these two stronger forms of bias-driven scepticism? Consider first strong bias-driven scepticism, which we saw was clearly present in Saul's work, and arguably also present in Alfano's work too. This sceptical line is effectively making it a requirement of knowledge

not only that one's true belief be the result of genuine cognitive ability, but also that one can rationally exclude the possibility that one's belief is subject to cognitive bias. Is it plausible that knowledge must satisfy this additional demand? At the very least it is a controversial claim to make, and will likely be rejected by many epistemologists of a broadly externalist persuasion (including proponents of the externalist wing of virtue epistemology, virtue reliabilism). After all, it is key to such proposals that a subject may have no rational basis at all for her beliefs. Just so long as her true beliefs are appropriately attributable to her exercise of reliable cognitive ability, then that can suffice to generate knowledge. Whether or not one is persuaded by the case for strong bias-driven scepticism will thus in part depend upon one's wider epistemological commitments. Moreover, if it turns out that the weaker intermediate bias-driven scepticism is problematic, then that will make it even harder to make a case for strong bias-driven scepticism.

What about intermediate bias-driven scepticism, which we take both Saul and Alfano to clearly endorse? Here too there are grounds for dispute. Consider a subject undertaking a complex piece of reasoning, and in the process forming a true belief. While there is plenty of genuine cognitive ability on display—careful attention to the relevant facts, accurate inference, and so on—there is also some cognitive bias in play, of which the subject is completely unaware (the use of the availability bias, say). Nonetheless, the cognitive bias does not prevent the subject from forming a true belief in the target proposition. Indeed, in order to keep the case as clean as possible, we can stipulate that the cognitive bias doesn't lead the subject to form any related false beliefs either. Here we have a case where the presence of the cognitive bias plays some explanatory role in the formation of the subject's true belief, but where the subject's cognitive abilities play a much larger explanatory role. If intermediate bias-driven scepticism is right, then this isn't a case of knowledge because the presence of the cognitive bias infects the epistemic standing in play. But, as with strong bias-driven scepticism, we seem to have a relatively austere epistemological thesis in play here, one that demands that knowledge is only *bona fide* if one's cognitive success is exclusively attributable to one's exercise of cognitive ability. And yet even proponents of the strongest form of virtue epistemology—known as *robust virtue epistemology*—argue that one's cognitive success only needs to be primarily attributable to one's exercise of cognitive agency, and hence are quite content to allow other factors to play a minor explanatory role.¹⁵ Robust virtue epistemology is thus compatible with the idea that there could be knowledge in the case just described.

Interestingly, one of the reasons why robust virtue epistemology is thought controversial also opens up the possibility that cognitive bias may play a much larger role in the formation of a true

belief without that belief thereby being prevented from counting as knowledge. In particular, robust virtue epistemology struggles to accommodate the *epistemic dependence* of knowledge, which is the way in which knowledge can be dependent upon factors outwith the subject's cognitive agency.¹⁶ Epistemic dependence comes in both *positive* and *negative* varieties. *Positive epistemic dependence* is when an agent manifests very little cognitive agency (i.e., much less than would normally suffice for knowledge), but where her cognitive success amounts to knowledge nonetheless because of factors external to her cognitive agency. *Negative epistemic dependence*, in contrast, is when an agent manifests a high level of cognitive agency (i.e., of a level that would ordinarily easily suffice for knowledge), but where the cognitive success does not amount to knowledge because of factors external to her cognitive agency.

A simple example of positive epistemic dependence involves testimonial knowledge acquired mostly via trust in an epistemically friendly environment (i.e., where testimony can generally be relied upon). You might, for instance, learn where the bus station is by simply asking someone you meet in the street. In doing so, one would ordinarily be exercising a certain level of cognitive agency—one wouldn't just ask anyone, for example, nor would one believe anything that one is told. One's cognitive agency is thus playing a significant explanatory role in one's cognitive success, and that is why one is in the market for testimonial knowledge (gullibility is not a route to testimonial knowledge). Crucially, however, in this case it is not one's own cognitive agency that is playing the overarching explanatory role in one's cognitive success, but rather one's informant. One's knowledge thus depends not only on one's cognitive performance, but also the cognitive performance of one's informant and the epistemically friendly nature of the environment.

The cleanest way to explain what negative epistemic dependence is by appeal to an epistemic twin earth case.¹⁷ Imagine that we have two counterpart agents, Sam and Sam*, where Sam is situated on earth and Sam* is on 'twin earth'. Stipulate that Sam and Sam* are microphysical duplicates with identical causal histories, and furthermore that they inhabit identical physical environments, both in terms of their *local* environment (i.e., their current environment which they are causally interacting with) and in terms of their *global* environment (i.e., the environment which they would be *normally* causally interacting with, where this could be different from their local environment). With these similarities in hand, let us stipulate now that Sam and Sam* both form the true belief that *p*, where the *only difference* in their circumstances concerns their respective *modal* environments—*viz.*, suppose there are close possible worlds where Sam* forms a false belief that *p* on the same basis as in the actual world, but there is no close possible world where Sam forms a false belief that *p* on the same basis as in the actual world.

A consequence of the difference in Sam and Sam*'s modal environments is that a true belief which is common to both of them can differ in terms of whether it is subject to knowledge-undermining *environmental epistemic luck*. After all, the true belief formed by Sam* on twin earth could easily have been false, though this is not so for Sam on earth. Crucially, however, Sam's and Sam*'s respective true beliefs do not differ in the extent to which they are attributable to their exercise of cognitive agency, even though their beliefs differ in the extent to which they are susceptible to knowledge-undermining environmental luck. What this means is that manifestations of cognitive agency that would ordinarily suffice for knowledge can nonetheless fail to suffice for knowledge, even when the target belief is true, thanks to factors outwith one's cognitive agency, such as what one's modal environment happens to be.¹⁸

The problem for robust virtue epistemology is that it cannot allow for cases of positive or negative epistemic dependence. In the former case, it is obliged to treat the subject as lacking knowledge, contrary to intuition. In the latter case, it is obliged to treat Sam*'s lucky true belief on twin earth as *bona fide* knowledge, again contrary to intuition. The import of this point for our discussion of bias-driven scepticism is that once one grants the phenomenon of epistemic dependence, then it is open to one to treat at least some cases involving cognitive bias as simply instances of positive epistemic dependence.

Consider again our agent from earlier who is undertaking a complex piece of reasoning, and in the process forms a true belief. The change we make to the story is to suppose that while there is a significant level of cognitive ability on display, there also a higher than before degree of cognitive bias on display too. For example, imagine that the explanatory burden when it comes to the subject's cognitive success is roughly equally split between the exercise of cognitive ability and the manifestation of cognitive bias. As before, however, we are stipulating here that the cognitive bias in play is not generating false beliefs at all, not just in the target proposition, but also in related propositions. Could a subject gain knowledge in such case, one where her success is not primarily attributable to her cognitive ability (though where cognitive ability is playing a large explanatory role)? Well, if one is willing to take the phenomenon of positive epistemic dependence seriously, then one should be at least sympathetic to this proposal, since absent the presentation of any relevant disanalogies, it looks to be very alike ordinary cases of positive epistemic dependence, where a subject just happens to be in the right conditions such that a limited display of cognitive ability can suffice for knowledge.

Our concern is not to press this point here. Rather, the claim we want to focus upon is that the plausibility of both intermediate and strong bias-driven scepticism seems to depend on further, and contentious, epistemological commitments, and for that reason both forms of scepticism are suspect. As we will see in the next section, however, there is a straightforward way of motivating analogous forms of scepticism, just so long as we shift our focus away from knowledge and towards understanding.

4. BIAS-DRIVEN SCEPTICISM, COGNITIVE ACHIEVEMENT AND UNDERSTANDING

In order to see why focussing on understanding rather than knowledge can help us resituate the stronger forms of bias-driven scepticism, we first need to make a few points about achievements, and cognitive achievements in particular. In general terms, an achievement is a success (within a given domain of endeavour) which is because of the manifestation of one's (relevant) ability. A familiar case used for illustrating this idea is archery. For instance, the archer's hitting the bull's-eye through a fortuitous gust of wind is clearly a success—the arrow after all has landed right where it is supposed to. But this success is not an achievement, as it is the gust of wind, rather than the archer's archery-relevant abilities, which primarily explains why the archer's shot hit the bull's-eye. In short, achievements require success and they require the relevant manifestation of ability, but they also require that the success be because of the ability, in the sense that it is the manifestation of ability that is the overarching element of a causal explanation of that success.

Interestingly, if robust virtue epistemology had been defensible, then it would have led to a conception of knowledge as cognitive achievement. That is, robust virtue epistemology is essentially the view that knowledge is cognitive success (i.e., true belief) that is because of (i.e., primarily attributable to) cognitive ability (i.e., epistemic virtue). The phenomenon of epistemic dependence reveals that this proposal is untenable, however, in that there are cognitive achievements that aren't knowledge (i.e., cases involving negative epistemic dependence), and instances of knowledge that aren't cognitive achievements (i.e., cases involving positive epistemic dependence). Knowledge and cognitive achievement thus come apart in both directions.

One putative advantage of thinking of knowledge as cognitive achievement is that because achievements are generally considered of special value, so we can explain why knowledge is of special

value.¹⁹ The idea that achievements have some special value is certainly plausible. For example, they seem to play an important role in a life of flourishing—there would be something lacking to a life that wasn't rich in achievements (which was regularly 'Gettierised', say). For the sake of argument, in what follows we will grant that achievements, and thus cognitive achievements are of special value (while being silent about what this value consists).²⁰

If cognitive achievements are of special value to us, then we should be alarmed if any argument comes along which shows that we have far fewer cognitive achievements than we hitherto supposed. With that in mind, consider again intermediate bias-driven scepticism. We saw above that this was problematic on account of the fact that knowledge allows for the possibility that one's cognitive success need not be exclusively attributable to one's exercise of cognitive agency. Indeed, if the argument for epistemic dependence is right, then one can sometimes have knowledge even where one's cognitive success is not even primarily creditable to one's exercise of cognitive agency. Interestingly, however, if cognitive achievements come apart from knowledge in the manner just described, then it follows that intermediate bias-driven scepticism, while problematic when applied to knowledge, would have straightforward application to cognitive achievement. In particular, while it is possible that some level of cognitive bias might be compatible with knowledge—on the grounds that it is a variety of positive epistemic dependence—that obviously won't work if we focus instead on cognitive achievement. After all, we granted above that the bias-driven sceptical argument does demonstrate that our cognitive successes in the relevant domains are not primarily creditable to our epistemic virtue, and that means that they are not candidates for cognitive achievements. The Saul-Alfano line thus does have the means to deprive us of an epistemic standing that is generally regarded as being of special value, by showing that we overestimate the role of our cognitive agency in our cognitive successes.

We can add an extra spin to this point by noting that the particular epistemic standing of *understanding-why* such-and-such is the case is plausibly a kind of cognitive achievement (henceforth, just 'understanding'). In particular, one of the present authors has previously argued that understanding-why comes apart from knowledge in just the same way that cognitive achievement does. In cases of negative positive dependence, where knowledge is lacking, the corresponding understanding is present. And in cases of positive epistemic dependence, in which the knowledge is present, the corresponding understanding is lacking.²¹

Rather than rehearsing these arguments here, we will instead focus on a particular example of positive epistemic dependence (which, as we saw above, is the kind of epistemic dependence of most

relevance to us), and show why knowledge, but not the corresponding understanding, is compatible with it. Let's take a case inspired by the particular kind of cognitive biases explored by Saul (2013a). Imagine that one is interviewing a range of candidates for a position and has formed the view that X is the best candidate, though in a way that was in part a result of cognitive ability and in part due to cognitive bias, albeit bias that in this case generates true rather than false beliefs. Suppose, for example, that stereotyping played at least some role in the formation of the belief, but in a fashion that neither led to a false belief in the target proposition, nor to false beliefs in the general vicinity of the target proposition.

As we saw above, on a plausibly formulated virtue-theoretic conception of knowledge, one which (unlike robust virtue epistemology) is compatible with epistemic dependence, one can contend that the target belief counts as genuine knowledge. That cognitive bias afflicts judgements of this kind means that one's epistemic virtues cannot plausibly bear the overarching explanatory load in one's cognitive success. But that is, as we've suggested in §3, compatible with them nonetheless playing a significant role, one that can sometimes be enough for knowledge.

Once one grants that the literature on cognitive bias downgrades the explanatory role that one's epistemic virtues can play in one's cognitive successes in these cases, however, then it immediately follows that this cognitive success doesn't amount to a cognitive achievement. But we contend that it also follows that one does not understand this candidate is the best candidate either, even though one knows this. This might initially seem surprising, but becomes more plausible once one considers some other cases where knowledge and understanding come apart in this.

A child, for example, can come to know that his family's house burned down because of faulty wiring because a parent told him so. But if he lacks a solid grasp of how faulty wiring could cause a fire, then he does not thereby understand why the house burned down. Or imagine the brilliant mathematician who can prove the theorem and the student who only believes it because the brilliant mathematician told him about it. While they both know that the proposition in question is true, only the mathematician understands why it is true. The crux of the matter is that when one has understanding, one has a kind of cognitive ownership of the fact in question—a concrete grasp of that fact, if you will. Note too that these are both cases of positive epistemic dependence, in that the cognitive success enjoyed by the subject with the knowledge but not the corresponding understanding is far less attributable to his cognitive agency, but is rather more dependent upon external factors. Cases of positive epistemic dependence show that one can gain knowledge without the kind of cognitive ownership characteristic of understanding, by in effect epistemically 'piggy-backing' off

factors which are external to one's cognitive agency. But the thought is that understanding is not like this.

Now consider again one's judgement that the candidate is the best person for the job. On the supposition that the individual's cognitive success is not primarily explained by her exercise of cognitive ability but also partly down to an unconscious bias, is one able to enjoy the kind of cognitive ownership of this fact that is characteristic of understanding? We suggest not. The positive epistemic dependence in play instead entails that one's epistemic position is more akin to the maths student who has to take the truth of the theorem on trust.

If the foregoing is right, then it's apparent how Saul and Alfano are in a position to recast intermediate bias-driven scepticism so that it targets understanding rather than knowledge. While the presence of cognitive bias in these cases can only undermine knowledge if one ignores the phenomenon of epistemic dependence, there is no such barrier to treating the cognitive bias in play as undermining understanding.

Interestingly, by shifting the target of bias-driven scepticism from knowledge to understanding, one can also give a fillip to strong bias-driven scepticism. For while it might be controversial to hold that knowledge always requires the kind of rational basis that is demanded by strong bias-driven scepticism, such that one must be able to rationally exclude the presence of cognitive bias, it is quite plausible that understanding might make such a demand. The point is that while knowledge is often understood along broadly externalist lines—such that one can know even while having very little in the way of reflective rational support, just so long as one's beliefs are in fact formed appropriately—it is far less plausible that understanding is susceptible to a comparable externalist treatment. To understand why such-and-such is the case is, after all, to have a reflective grip on the matter. How could one possess such a reflective grip while lacking a rational basis of just this sort?

Consider, for example, a counterpart to the person in the example just cited, who has made the correct judgment about whom to hire. Imagine now, however, that while she has formed her belief in a way unaffected by cognitive bias, she could very easily have formed her belief in part or in whole due to cognitive bias and have been none the wiser. Since she has no reflective basis for excluding such a live possibility, does she really understand why this candidate is right one for the job? We suggest not. Like cognitive achievement, understanding requires a particular kind of cognitive ownership that isn't essential for merely knowing. But, unlike cognitive achievement, that cognitive ownership also essentially involves a reflective grip on the matter in hand, one that would require the

subject to have a rational basis to exclude live error-possibilities. Strong bias-driven scepticism is thus back up-and-running, albeit now targeted at understanding rather than knowledge.

In summary, although the kind of bias-driven scepticism presented (in different forms) by Alfano and Saul is problematic when applied to knowledge, because it depends on wider epistemological commitments, those difficulties disappear once we refocus this challenge specifically on the distinct epistemic standing of understanding.²²

REFERENCES

- Alfano, M. (2012). 'Expanding the Situationist Challenge to Responsibilist Virtue Epistemology', *Philosophical Quarterly* 62, 223-49.
- (2014). 'Expanding the Situationist Challenge to Reliabilism About Inference', *Virtue Epistemology Naturalized*, (ed.) A. Fairweather, 103-22, Dordrecht, Holland: Springer.
- Axtell, G. (1997). 'Recent Work on Virtue Epistemology', *American Philosophical Quarterly* 34, 1-27.
- Beulac, G., & Kenyon, K. (2014). 'Critical Thinking Education and Debiasing', *Informal Logic* 34, 341-63.
- Bertrand, M., & Mullainathan, S. (2003). 'Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination', *American Economic Review* 94: 991-1013.
- Carter, J. A. (2015). *Metaepistemology and Relativism*, London: Palgrave Macmillan.
- Carter, J. A., & Kallestrup, J. (2015). 'Extended Cognition and Propositional Memory', *Philosophy and Phenomenological Research* (Online First, DOI: DOI: 10.1111/phpr.12157).
- Carter, J. A., & Pritchard, D. H. (2015). 'Epistemic Situationism, Epistemic Dependence and the Epistemology of Education', *Epistemic Situationism*, XX-XX, (eds.) M. Alfano & A. Fairweather, Oxford: Oxford University Press.
- Correll, J., Urland, G. R., & Ito, T. A. (2006). 'Event-Related Potentials and the Decision to Shoot: The Role of Threat Perception and Cognitive Control', *Journal of Experimental Social Psychology* 42, 120-28.
- deTurck, M. A., Harszslak, J. J., Bodhorn, D. J., & Texter, L. A. (1990). 'The Effects of Training Social Perceivers to Detect Deception from Behavioral Cues', *Communication Quarterly* 38, 189-99.
- Doris, J. (1998). 'Persons, Situations, and Virtue Ethics', *Noûs* 32, 504-40.
- (2002). *Lack of Character: Personality and Moral Behavior*, Cambridge: Cambridge University Press.
- Harman, G. (1999). 'Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error', *Proceedings of the Aristotelian Society* 99, 316-31.
- (2000). 'The Nonexistence of Character Traits', *Proceedings of the Aristotelian Society* 100, 223-26.
- (2003). 'No Character or Personality', *Business Ethics Quarterly* 13, 87-94.
- Kruger, J., & Dunning, D. (1999). 'Unskilled and Unaware of it: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments', *Journal of Personality and Social Psychology* 77, 1121-34.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*, Oxford: Oxford University Press.
- Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*, Cambridge: Cambridge University Press.
- (2013). 'A (Different) Virtue Epistemology', *Philosophy and Phenomenological Research* 85, 1-26.
- Greenwald, A. G., Oakes, M. A., & Hoffman, H. G. (2003). 'Targets of Discrimination: Effects of Race on Responses to Weapons Holders', *Journal of Experimental Social Psychology* 39, 399-405.
- Holroyd, J., & Sweetman, J. (2015). 'The Heterogeneity of Implicit Bias', *Implicit Bias and Philosophy*, (ed.) M. Brownstein & J. Saul, XX-XX, Oxford: Oxford University Press.
- Kallestrup, J., & Pritchard, D. H. (2012). 'Robust Virtue Epistemology and Epistemic Anti-Individualism', *Pacific Philosophical Quarterly* 93, 84-103.
- (2013). 'Robust Virtue Epistemology and Epistemic Dependence' (with J. Kallestrup), *Knowledge, Virtue and Action: Putting Epistemic Virtues to Work*, (eds.) T. Henning & D. Schweikard, ch. 11, (Routledge, 2013).
- (2014). 'Virtue Epistemology and Epistemic Twin Earth', *European Journal of Philosophy* 22, 35-357.

- Kvanvig, J. L. (2003). *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.
- Moss-Racusin, C. A., Dovidio, J. F., Brescoll, V. L., Graham, M. J., & Handelsman, J. (2012). 'Science Faculty's Subtle Gender Biases Favor Male Students', *Proceedings of the National Academy of Sciences* 109, 16474-79.
- Peters, D. P., & Ceci, S. G. (1982). 'Peer-Review Practices of Psychological Journals: The Fate of Published Articles, Submitted Again', *Behavioral and Brain Sciences* 5, 187-95.
- Pritchard, D. H. (2002). 'Recent Work on Radical Skepticism', *American Philosophical Quarterly* 39, 215-57.
- (2008). 'Radical Scepticism, Epistemic Luck, and Epistemic Value', *Proceedings of the Aristotelian Society (suppl. vol.)*, 82, 19-41.
- (2009). 'Knowledge, Understanding and Epistemic Value', *Epistemology (Royal Institute of Philosophy Lectures)*, (ed.) A. O'Hear, 19-43, Cambridge: Cambridge University Press.
- (2010). 'Achievements, Luck and Value', *Think* 25 (2010), 1-12.
- (2012). 'Anti-Luck Virtue Epistemology', *Journal of Philosophy* 109, 247-79.
- (2014a). 'Knowledge and Understanding', *Virtue Scientia: Bridges Between Virtue Epistemology and Philosophy of Science*, (ed.) A. Fairweather, 315-28, Dordrecht, Holland: Springer.
- (2014b). 'Re-evaluating the Epistemic Situationist Challenge to Virtue Epistemology', *Naturalizing Epistemic Virtue*, (eds.) A. Fairweather & O. Flanagan, 143-54, Cambridge: Cambridge University Press.
- (2014c). 'The Modal Account of Luck', *Metaphilosophy* 45, 594-619.
- (2015a). *Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing*, Princeton, NJ: Princeton University Press.
- (2015b). 'Anti-Luck Epistemology and the Gettier Problem', *Philosophical Studies* 172, 93-111.
- (2015c). 'Risk', *Metaphilosophy* 46, XX-XX.
- Pritchard, D. H., Millar, A., & Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations*, Oxford: Oxford University Press.
- Pritchard, D. H., & Turri, J. (2011). 'The Value of Knowledge', *Stanford Encyclopedia of Philosophy*, (ed.) E. Zalta, <http://plato.stanford.edu/entries/knowledge-value/>.
- Saul, J. (2013a). 'Scepticism and Implicit Bias', *Disputatio* 5, 243-63.
- (2013b). 'Implicit Bias, Stereotype Threat and Women in Philosophy', *Women in Philosophy: What Needs to Change?*, (eds.) F. Jenkins & K. Hutchison, ch. 2, Oxford: Oxford University Press.
- Slovic, P. (1987). 'Perception of Risk', *Science* 236, 280-85.
- Sosa, E. (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.
- Strack, F., & Mussweiler, T. (1997). 'Explaining the Enigmatic Anchoring Effect: Mechanisms of Selective Accessibility', *Journal of Personality and Social Psychology* 73, 437.
- Tversky, A., & Kahneman, D. (1973). 'Availability: A Heuristic for Judging Frequency and Probability', *Cognitive Psychology* 5, 207-32.
- (1974). 'Judgment Under Uncertainty: Heuristics and Biases', *Science* 185, 1124-31.
- Unkelbach, C., Forgas, J. P., & Denson, T. F. (2008). 'The Turban Effect: The Influence of Muslim Headgear and Induced Affect on Aggressive Responses in the Shooter Bias Paradigm', *Journal of Experimental Social Psychology* 44, 1409-13.

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

NOTES

¹ On reflection, that local varieties of scepticism could be intellectually devastating should not be surprising because—as with any assessment of costs—we must take in to account when assessing sceptical costs not just the quantity, but also the *quality* of what is being denied. Thus, a local form of radical scepticism could be more devastating than a radical counterpart in virtue of targeting a more valuable epistemic standing. For more on the relationship between debates about epistemic value and the problem of radical scepticism, see Pritchard (2008).

² For further discussion of the nature of sceptical challenges, see Pritchard (2002; 2015*a*, part one).

³ See Beaulac & Kenyon (2014, 344). They note that the characterisation offered ‘accommodates the idea that biases can skew a process in a way that makes its outcome inaccurate or otherwise wrong, but it also leaves open the prospect that biases play a role in truth-conducive reasoning processes and morally unproblematic judgments or attitudes.’

⁴ In other work, Saul (2013*b*) uses ‘implicit bias’ more inclusively, as ‘unconscious biases that affect the way we perceive, evaluate, or *interact* with people from the groups that our biases ‘target’ (*our italics*). The broader classification allows for implicit biases as such to include biases that terminate in *action*, rather than merely representation. In focusing on *cognitive* biases, our focus will be on implicit biases that terminate in representation, which are exactly the kinds of implicit biases that Saul (2013*a*) is primarily appealing to. For a comprehensive treatment of implicit bias, see Holroyd & Sweetman (2015).

⁵ See Moss-Racusin *et al* (2012) and Bertrand & Mullainathan (2003) for some example studies.

⁶ See Correll, Urland & Ito (2006), Greenwald, Oakes & Hoffman (2003), and Unkelbach, Forgas & Denson (2008). See also Saul (2013*a*, 245-46). Saul also discusses the related phenomenon of ‘prestige bias’, in which papers *already published* in certain high-profile journals were resubmitted with false names and non-prestigious affiliations. However, as Saul notes:

“Only 8% detected that the papers had already been submitted, and 89% were rejected, citing serious methodological errors (and not the one they should have cited—plagiarism). This makes it clear that institutional affiliation has a dramatic effect on the judgments made by reviewers (either positively, negatively, or both).” (Saul 2013*a*, 245)

See Peters & Ceci (1982) for the seminal study of this effect.

⁷ See Saul (2013*a*, 245-46). Implicit biases culminating in inaccurate judgments of the testimony of individuals from stigmatised groups, a phenomenon known as testimonial injustice, is explored most notably by Fricker (2007).

⁸ For more on the distinction between virtue responsibility and virtue reliabilism, see Axtell (1997).

⁹ See, for example, Slovic (1987). For a discussion of the availability bias at work in responses to the 2014 Ebola outbreak, see <http://www.psychologytoday.com/blog/everybody-is-stupid-except-you/201412/the-heuristic-caused-the-ebola-panic-2014>. Note that risk perceptions are also subject to a range of other well-studied biases, including the anchoring-bias (e.g., probability estimates for a range of risks are typically within a range that tracks the first risk probability made salient). More generally, the anchoring bias involves interpreting a range of information in terms of the first piece of information—*viz.*, the ‘anchor’. See Strack & Mussweiler (1997). For more on the empirical work on risk, and some of the philosophical issues it raises, see Pritchard (2014*c*; 2015*e*).

¹⁰ To think otherwise is to succumb to the conjunctive fallacy—*viz.*, the probability of a conjunction cannot be higher than the probability of one of its conjuncts.

¹¹ In fact, there is some cause to doubt whether the subjects in this case are making the conjunctive fallacy, but in keeping with the general spirit of the paper we will take the empirical work cited (and what it purports to show) at face value. For further critical discussion of this case, see Pritchard (2014*b*).

¹² Note that our focus here is on negative cognitive biases, which is those biases that undermine the reliability of one’s cognitive powers. Some cognitive biases are, in fact, positive, in that they can aid the reliability of our cognitive powers, and this can complicate the epistemic situationist case against virtue epistemology. For further discussion of this point, see Pritchard (2014*b*) and Carter & Pritchard (2015).

¹³ In fact, the present authors have argued elsewhere that Alfano has almost certainly underplayed his hand here, in that the arguments he offers against inferential knowledge apply to lots of non-inferential knowledge too, even on a virtue reliabilist construal. See Carter & Pritchard (2015). We will set this one side in what follows.

¹⁴ See Carter (2015, ch. 4) for further discussion of the arguments put forward by Saul and Alfano.

¹⁵ Versions of this thesis have been defended by Zagzebski (1996), Sosa (2007; 2009), and Greco (2009).

¹⁶ For the main discussions of epistemic dependence, see Kallestrup & Pritchard (2012; 2013; 2014). Note that this notion is rooted in the critique of robust virtue epistemology offered in Pritchard (e.g., 2009; 2012) and Pritchard, Millar & Haddock (2010, chs. 1-4).

¹⁷ See Kallestrup & Pritchard (2012) for the original description of the epistemic twin earth case.

¹⁸ For more on environmental epistemic luck, and how it differs from other kinds of knowledge-undermining luck, see Pritchard (2009; 2012; 2015*e*) and Pritchard, Millar & Haddock (2010, chs. 1-4).

¹⁹ As Greco puts it:

“Knowledge is an achievement in a sense that lucky guesses (and the like) are not. This is an improvement over accounts that make the value of knowledge merely practical or instrumental. Plausibly, we value knowledge (as we value achievement in general) “for itself,” over and above its practical or instrumental value.” (Greco 2013, 2)

²⁰ Pritchard has explored the value of achievements, and of cognitive achievements in particular, in a number of places—see, especially, Pritchard (2009; 2010) and Pritchard, Millar & Haddock (2010, chs. 1-4).

²¹ See Pritchard (2009; 2014*a*) and Pritchard, Millar & Haddock (2010, chs. 1-4) for the development of the view that knowledge and understanding-why come apart in this manner, and hence that understanding-why is a kind of cognitive achievement.

²² Acknowledgements.