

UC Irvine

UC Irvine Previously Published Works

Title

The Epistemology of Cognitive Enhancement

Permalink

<https://escholarship.org/uc/item/2d45t6qv>

Journal

The Journal of Medicine and Philosophy A Forum for Bioethics and Philosophy of Medicine, 44(2)

ISSN

0360-5310

Authors

Carter, J Adam

Pritchard, Duncan

Publication Date

2019-03-16

DOI

10.1093/jmp/jhy040

Copyright Information

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

Peer reviewed

THE EPISTEMOLOGY OF COGNITIVE ENHANCEMENT

J. Adam Carter & Duncan Pritchard

University of Edinburgh

ABSTRACT. A common epistemological assumption in contemporary bioethics held by both proponents and critics of non-traditional forms of cognitive enhancement is that cognitive enhancement aims at the facilitation of the accumulation of human knowledge. This paper does three central things. First, drawing from recent work in epistemology, a rival account of cognitive enhancement, framed in terms of the notion of *cognitive achievement* rather than knowledge, is proposed. Second, we outline and respond to an axiological objection to our proposal that draws from recent work by Leon Kass (2004), Michael Sandel (2009), and John Harris (2011) to the effect that ‘enhanced’ cognitive achievements are (by effectively removing obstacles to success) not worthy of pursuit, or are otherwise ‘trivial’. Third, we show how the cognitive achievement account of cognitive enhancement proposed here fits snugly with recent active externalist approaches (e.g., extended cognition) in the philosophy of mind and cognitive science.

Keywords: *cognitive enhancement; bioethics; epistemic value; cognitive achievement; knowledge*

1. COGNITIVE ENHANCEMENT AND KNOWLEDGE

One of the most provocative and pressing topics in recent bioethics concerns human *enhancement*.¹ The latest science and medicine makes it increasingly possible to improve human functioning along several different dimensions: physical, cognitive and (arguably) even moral.² On the assumption that such improvements are available, there is a *pro tanto* instrumental reason to pursue them. After all, to the extent that we are able to function more effectively by being enhanced, we can better achieve our adopted ends. But as critics of human enhancement suggest, the overall balance of reasons may nonetheless speak against human enhancement. That is, all things considered, it may be that we’re better off *not* improving certain dimensions of our functioning, at least in certain ways.

The variety of improvements to human functioning that will be of interest in what follows will be specifically *cognitive* improvements—*viz.*, improvements in aspects of human intelligence, such as memory, reasoning, computation and decision making,³ as opposed to (for example) improvements in the affective (i.e., emotional) or conative (volitional) states or

dispositions. Obviously, not all kinds of cognitive improvements are controversial. Accordingly, bioethicists typically distinguish between the following two kinds of cognitive improvements:

Therapeutic Cognitive Improvements: These are aimed at *correcting* a specific pathology or cognitive defect (e.g., treatments aimed at slowing the progression of Alzheimer's).

Cognitive Enhancements: These are (roughly) interventions that improve cognitive functioning *non-therapeutically* (e.g., by ways that do not involve the repairing of a cognitive defect or dysfunction).⁴

In the latter case, the aim of the improvement is not restoration to healthy cognitive functioning, but rather to (for instance) gain some kind of cognitive advantage.⁵

A further refinement is now needed, as not all attempts to gain a cognitive advantage are controversial.⁶ One very natural (and uncontroversial) way to gain a cognitive advantage relative to one's normal level of cognitive functioning is through entirely traditional methods, such as studying.⁷ While traditional cognitive enhancement has a long history, it has only relatively recently become possible to improve human capacities in non-traditional ways—e.g., by making use of the latest medicine and science to improve cognitive capacities in healthy individuals, through such methods as nootropics or 'smart drugs' (e.g., Adderall, Ritalin, Provigil, Oxiracetan), implants (e.g., neuroprosthetics⁸), direct brain-computer interfaces, and genetic engineering.⁹

One thriving research area surrounds the matter of just what distinguishes traditional cognitive enhancement (such as studying) from non-traditional cognitive enhancement.¹⁰ We will not be engaging with this issue here.¹¹ Rather, we want to highlight a very simple—and often overlooked—background assumption which is widely endorsed on both sides of debates about the moral permissibility of (non-traditional) cognitive enhancement. Call this the *knowledge account of cognitive enhancement*:

Knowledge Account of Cognitive Enhancement (KA_{CE}): Cognitive enhancement aims at the facilitation of the accumulation of knowledge.

Explicit, and implicit, commitments to this proposal are widespread. Consider, for instance, Michael Hauskeller's remark that:

Cognitive enhancements are, roughly put, all interventions that, through the manipulation of the human brain, improve the *human knowledge situation* by facilitating or accelerating knowledge acquisition, processing, storage, application or range. (Hauskeller 2013, 117)

Likewise, Ingmar Persson and Julian Savulescu (2008, 3)—two of the most influential contemporary bioethicists—describe cognitive enhancement as generative of a 'faster increase of

knowledge.’ Moreover, they explicitly regard cognitive enhancement as being geared to ‘speed up the *acquisition of knowledge*.’¹² (Persson & Savulescu 2008, 7)

Two points of clarification will be helpful here. First, ‘enhancement’ is a normative concept, and one that makes sense to apply only with some dimension of evaluation in mind. A cognitive enhancement, as such is—to a first approximation—an improvement from the point of view where what matters is some [to be specified] aim of cognition, and the knowledge account fills in this aim with ‘knowledge’. For philosophers who deny that cognition has any such aim at all, a presupposition of this way of thinking of cognitive enhancement will of course be rejected.¹³ However, given the thought that there are certain values or aims, such as knowledge, which *govern* the practice of cognizing¹⁴ (and structure our evaluations within this practice), the notion of cognitive enhancement will be tied to the matter of how such a governing value—or *aim*—is to be specified. In epistemology, at least, the matter of how to specify such an aim is an increasingly contested issue.¹⁵ ‘Knowledge’ is a popular position, but the value of knowledge (in comparison with other rival epistemic standings) has increasingly been called into question—a point we’ll return to.

Second, it’s important to note that KA_{CE} is compatible with the thought that cognitive enhancement can and often will involve improvements of *abilities* (as opposed to simply direct improvements to one’s knowledge states). What KA_{CE} says is that what *makes* an improvement of human ability an instance of *cognitive* enhancement (as opposed to, say, moral enhancement) is that the agent becomes (*via* the improvement) better situated with respect to the aim of acquiring knowledge.

2. CONTRA THE KNOWLEDGE ACCOUNT OF COGNITIVE ENHANCEMENT

KA_{CE} is initially very plausible. Compare, for instance, KA_{CE} with a much weaker (and less plausible) account of cognitive enhancement in terms of true belief:

True Belief Account of Cognitive Enhancement (TB_{CE}): Cognitive enhancement aims at the facilitation of the accumulation of true belief.

To make concrete just why TB_{CE} runs in to problems, just consider the following ‘meta-incoherence’-style case, where (in short) one gets, thanks to an implant, reliably true beliefs, though the agent has no conception of the source of the reliability:

Chip: Chip has (unbeknownst to him) undergone an experimental procedure in which a complex computer chip has been implanted in the parietal lobe of his cerebral cortex. Whenever Chip

considers which direction he is currently facing, he compulsively forms a correct belief. Whilst Chip's directional orientation beliefs are invariably correct he has no idea why.¹⁶

After the procedure, Chip acquires more true beliefs about his current directional orientation than he would have had without the implant. However, Chip's 'guaranteed-true' directional beliefs are epistemically defective in an important respect.¹⁷ This is that there is a lack of coherence between, as Jonathan Kvanvig (2003, 454) has put it, 'the understanding [*Chip*] has of the reliability of various ways of forming and holding [*the relevant*] beliefs and how the beliefs are formed in these cases.' In short, Chip is epistemically better off than he was before, but not nearly as well off as he could be—i.e., not as well off as he would be had this incoherence been overcome.

To bring into focus exactly why the case of Chip can be converted into a problem for TB_{CE}, just contrast the case of *Chip* with a parallel case of *Chip** for whom we stipulate that the implant results in equally reliable directional-orientation beliefs though (thanks to more nuanced integration) with *no* residual meta-incoherence.¹⁸ On a continuum of cognitive enhancement, *Chip** is surely further along the continuum in the direction of the cognitive ideal than *Chip*. But TB_{CE} lacks the resources to explain why this could ever be so, as both *Chip** and *Chip*, thanks to their respective implanted chips, accumulate the same gain in true beliefs relative to their normal (healthy) baseline.

KA_{CE} by contrast has an elegant explanation for why *Chip** is further along the cognitive enhancement continuum than *Chip*. The reasoning is simple. First, cognitive enhancement (according to KA_{CE}) *aims* at the facilitation of the accumulation of knowledge. Second, given that the reliability at play in meta-incoherence cases (like the case of *Chip*) is widely regarded as insufficient to support knowledge,¹⁹ *Chip**'s implant better achieves this aim—*viz.*, the aim of facilitating the accumulation of knowledge—than does *Chip*'s.²⁰

Though KA_{CE} looks very plausible in comparison with TB_{CE}, we want to suggest now that KA_{CE} is ultimately unsatisfactory as a thesis about cognitive enhancement. To this end it will be helpful to consider a different sort of case.

Lestrade. Lestrade, a normal healthy adult improves (through, say, genetic enhancement) certain cognitive abilities that enable him to much better than previously (e.g., at normal levels of functioning) grasp explanatory connections within a given body of information. Lestrade is assisting the police with a murder investigation and is given a list of 10 facts about the crime scene, and a list of three leading suspects (Mr. X, Mr. Y, and Mr. Z). He easily (and correctly) works out very quickly that Mr. X must be the murderer. Prior to the cognitive enhancement, Lestrade would not have been able to work this out (at least not without great difficulty anyway).²¹

According to KA_{CE} , the genetic modification underwent by Lestrade is an instance of cognitive enhancement in virtue of better situating Lestrade with respect to the aim of knowledge acquisition. That Lestrade comes to know something via this cognitive enhancement that he wouldn't have otherwise known thus accords with KA_{CE} . What we want to draw attention to, however, is that Lestrade's enhanced abilities also facilitate his ability to gain a different kind of epistemic standing. This is *understanding-why*—i.e., the state that one attains when one understands *why* something is the case (e.g., why the television isn't working).

For ease of expression, we will talk of understanding-why and understanding interchangeably. Note, however, that in the contemporary literature the former is regarded as a sub-species of the latter, in that there are other kinds of understanding, such as what is known as *holistic understanding* (as when one understand a subject matter, like quantum physics) that have different properties. Nonetheless, understanding-why is at least treated as a paradigm kind of understanding, and hence this simplification ought to be harmless for our purposes.²²

Notice that, as we have described the case, Lestrade doesn't merely know that the murderer must be Mr. X, but also understands why this is so, in virtue of being able to reason his way to this conclusion. In order to see this point, compare Lestrade with a colleague who lacks Lestrade's insight into this case and who merely knows that it must be Mr. X that is the murder because Lestrade—recognised now as the expert in such matters—has testified to this fact. Lestrade's colleague has knowledge, but unlike Lestrade he lacks the corresponding understanding of why the murderer must be Mr. X. If he is asked why he thinks that it's Mr. X, the best he can do is cite Lestrade's testimony. Lestrade, in contrast, can cite the pattern of reasoning that led to this conclusion, and in doing so manifest his understanding of why Mr. X must be the murderer. Knowing that something is the case thus does not entail understanding why it must be case.

Obviously, this point is by the by—for the purposes of assessing KA_{CE} —if understanding why something is so just *is* to have a relevant piece of (propositional) knowledge (i.e., if understanding just is a kind of knowing, which is of course compatible with knowledge not entailing understanding). If so, pointing out that Lestrade's enhancement facilitates understanding stands in no obvious tension with KA_{CE} . And indeed, the thesis that understanding-why is just a matter of having a certain kind of knowledge—typically, knowing the relevant causes of the event in question—has had a number of influential takers, especially at the intersection of epistemology and the philosophy of science.²³

However, the tide has been turning against such reductive 'knowledge' accounts of understanding.²⁴ The problem with these proposals is that they are unable to capture the

particular manner in which an agent can have understanding while lacking the corresponding knowledge. The issue here concerns epistemic luck, and a crucial distinction between two kinds of knowledge-undermining epistemic luck. In general, it is widely agreed that knowledge excludes high levels of *veritic epistemic luck*, where this means that when one knows, one could not (by forming a belief on the same basis) easily be wrong. This is often given a modal spin and characterised as the thesis that knowledge demands *safety*, where this means that one's belief-forming method would not have easily resulted in error in close possible worlds.²⁵ The language of 'safety' is instructive—what we want from a theory of knowledge is an account that avoids high levels of *epistemic risk*, and any belief that is unsafe in this sense would be epistemically very risky indeed.²⁶

Crucially, however, veritic epistemic luck comes in two forms: *intervening epistemic luck* and *environmental epistemic luck*.²⁷ The former is the standard kind of knowledge-undermining luck at issue in typical Gettier-style cases, as when one believes that there is a sheep in the field because one is looking at a sheep-shaped object (which, unbeknownst to one, is hiding the genuine sheep in the field behind). The latter is less common. It concerns cases where there is none of the disconnect between belief and fact that one finds in standard Gettier-style cases, but where, nonetheless, there is some feature of the environment that ensures that the belief is unsafe regardless (and hence that the level of epistemic risk is unduly high). Barn façade cases are an obvious case in point. After all, one really does see a genuine barn in such scenarios, in conditions which are generally optimal for barn-perception. And yet, intuitively, one does not know in such cases because one is in an environment which has features (on account of the widespread barn façades) that ensures that the true belief, so formed, is unsafe.

This distinction is important to the debate about knowledge and understanding because only the former is incompatible with both kinds of epistemic luck. Understanding-why, in contrast, while incompatible with intervening epistemic luck, is entirely compatible with environmental epistemic luck. One can illustrate this point by comparing the following two cases, where the first involves intervening epistemic luck and the second environmental epistemic luck:

Alexander

Alexander comes home to find his house in flames. He approaches someone who looks as if she is the fire officer in charge and asks her what the reason for the fire is. He is told by this person that the reason why his house is burning down is faulty wiring, and this coheres with his wider set of beliefs (e.g., about how faulty wiring can cause a house fire). The person Alexander is speaking to, however, is not in fact a fire officer at all but instead someone who is merely dressed in a fire officer's uniform and who is on her way to a fancy dress party. Even so, what she said to Alexander was true: faulty wiring *was* the cause of this fire.

Ethan

Ethan comes home to find his house in flames. He approaches someone who looks as if she is the fire officer in charge and asks her what the reason for the fire is. He is told by this person that the reason why his house is burning down is faulty wiring, and this coheres with his wider set of beliefs (e.g., about how faulty wiring can cause a house fire). The person Ethan is speaking to is indeed the fire officer in charge of this blaze, and what she says about the cause of the fire is both true and appropriately epistemically grounded (i.e., the fire officer *knows* that this is the cause of the fire). Crucially, however, all the other people in the vicinity who are dressed as fire officers are in fact on their way to a fancy dress party and have nothing to do with the fire brigade. Had Ethan asked one of them what the cause of the fire was, however, then they would have kept this fact from him and he would have believed their testimony.²⁸

What both cases share is that the subject gains a true belief, albeit one which is unsafe—i.e., given how it was formed, it could very easily have been a false belief. Accordingly, insofar as knowledge demands safety, as is typically supposed, then neither agent gains knowledge. Notice, however, what is different about the two cases. Whereas Alexander gains his true belief about the cause of the fire from a fake fire officer who is making his response up, Ethan gains his true belief by talking to a genuine knowledgeable fire officer (albeit one that he wouldn't have been able to distinguish from the fake fire officers in the vicinity). I doubt anyone would be inclined to claim that one can gain an understanding of why one's house burned down by believing the word of someone who is making up his explanation, so Alexander lacks understanding-why as well as knowledge. But given that Ethan has a grasp of how faulty wiring could be the cause of a house fire, what would prevent us from crediting him with an understanding of why his house burned down? After all, his belief in this regard, while unsafe, is acquired from a knowledgeable source and coheres with his wider set of beliefs.

If the foregoing is correct, then it follows that agents can have understanding-why without the corresponding knowledge, and hence that a knowledge account of understanding is in trouble (i.e., the view that all understanding-why is in virtue of the possession of propositional knowledge). In particular, it doesn't seem at all plausible that Ethan's possession of understanding-why is in virtue of what he knows. The immediate upshot of this point, however, is that insofar as we follow standard practice by treating understanding as a distinctively valuable kind of epistemic status, then there is a standing problem for the idea that one can accommodate cases of cognitive enhancement via KA_{CE} .²⁹ In particular, it seems that there can be cases of cognitive enhancement which won't be captured by KA_{CE} because the enhancement generates understanding-why but not the corresponding knowledge. Indeed, just imagine a version the Lestrade case where there is environmental epistemic luck in play. Lestrade is now unable to gain knowledge of the target proposition. But the cognitive enhancement involved in his heightened cognitive abilities is still on display, and generating understanding-why, albeit understanding-why which is not in virtue of his possession of knowledge.³⁰

Before moving on, though, we want to quickly note a potential line of objection. This proceeds by noting that, *even if*, as we've argued, understanding-why is not a matter of knowledge-that, so long as understanding-why is equivalent to, or grounded in, knowing-*why* as opposed to knowing that, then the foregoing discussion poses no insuperable problem for KA_{CE}. After all, if this were true, then it would follow that understanding-why can ultimately be accounted for in terms of a *kind* of knowledge—i.e., knowledge-why—even if not knowledge-that.

The problem with this anticipated line of reasoning is that, at least according to the prevailing view in mainstream epistemology, knowledge-why, like other kinds of knowledge-wh constructions (such as knowledge-where, knowledge-what, and so on), just is a *kind* of propositional knowledge. For example, as Jason Stanley (2011) puts it, drawing from the standard semantics for knowledge-wh constructions (e.g., Karttunen 1977), if Diana knows why the park was closed, then this is in virtue of Diana knowing, of some reason *r*, *that r* is the reason that the park is closed. And this is just propositional knowledge. Accordingly, then, understanding-why stands in relation to knowledge-why in the same way that understanding-why stands in relation to knowledge-that. And thus the foregoing problems for KA_{CE}, remain.³¹

3. THE COGNITIVE ACHIEVEMENT ACCOUNT OF COGNITIVE ENHANCEMENT

It looks very plausible then that understanding affords one a distinctive kind of cognitive good that one doesn't automatically possess in virtue of simply knowing facts. And note that this is so even if the accumulation of knowledge generally leaves one cognitively better off than the accumulation of baseless or irrational true beliefs. It follows that KA_{CE} faces serious problems. Thus far, our case against KA_{CE} has been negative. Let us now take a more positive tack and try to motivate an alternative account which is more promising.

To begin with, it will be worth reflecting on what the relationship is between knowledge and understanding-why, insofar as a knowledge account of understanding-why is rejected. We saw in the last section that cases of environmental epistemic luck demonstrate that one can have understanding-why without the relevant knowledge. It's also worth noting—though this point isn't strictly incompatible with the knowledge account of understanding-why—that one can have knowledge without the corresponding understanding-why. Imagine, for example, that a fire officer tells an 8-year-old child that the school burned down because of faulty wiring, and he

believes the fire officer on this basis. Nonetheless, although he understands the words in play (and hence can form the relevant belief), he might have no real grasp of how it could be that faulty wiring could lead to such an event. In such a case he would know the target proposition while lacking the corresponding understanding-why.

Knowledge and understanding-why thus come apart in both directions. Moreover, when they do come apart, the intuition seems to be that understanding enjoys a kind of epistemic value that mere knowledge (in the absence of understanding) lacks. After all, it is better to not only know that the house burned down because of faulty wiring but also to understand why this is the case. Or consider cases of environmental epistemic luck where understanding-why is present but knowledge is lacking. Doesn't the subject in such cases enjoy an epistemic standing which is distinctively valuable, even despite the lack of knowledge?

There is a neat explanation available of why understanding-why comes apart from knowledge in this fashion, and which also explains why it is distinctively valuable. This is that understanding is a particular kind of cognitive achievement (call this the *cognitive achievement account of understanding*).³² In order to see the attraction of this view, consider first the nature of achievements more generally. On the standard view of achievements in epistemology—one that has been articulated by (among others) Ernest Sosa (e.g., 1991; 2007; 2009; 2011; 2015), John Greco³³ (2003; 2010; 2012), John Turri (*forthcoming*), Linda Zagzebski (1996) and Pritchard (2009*a*; 2009*b*; 2009*c*; 2010*a*; 2012)—achievements are successes that are primarily attributable or creditable to one's exercise of relevant ability. So for an archer to exhibit the achievement of hitting the target, it's not enough that she's successful, as that could be down to dumb luck. Moreover, it's not even enough that there is the combination of success and the manifestation of relevant ability, since one can always 'Gettierize' the relationship between the two by bringing in intervening luck. A skilled archer who fires her bolt at the target, but who only succeeds in hitting the target because a dog snatches it in mid-flight and deposits it in the bulls-eye does not exhibit an achievement. For while there is both success and the manifestation of the relevant ability on display, the former is not primarily attributable to the latter (being attributable, instead, to the intervention of the dog).

Interestingly, however, while achievements are incompatible with intervening luck, they are entirely compatible with environmental luck. Imagine now that the very same archer skilfully fires her bolt and thereby hits the target, where something very nearly intervened but didn't (e.g., the dog jumps to snatch the arrow, but narrowly misses). As before, we have a success which is unsafe (i.e., which could so very easily have been a failure). But the environmental luck in play

does not prevent this success from being an achievement. After all, given that nothing did in fact intervene, the archer's success *is* primarily attributable to her relevant abilities.

What goes for achievement goes for cognitive achievement too—i.e., cognitive success (e.g., true belief³⁴) that is primarily attributable to the manifestation of relevant cognitive ability, where the notion of ‘cognitive ability’ or intellectual virtue at play here is ordinarily conceived of as reliable, appropriately integrated traits of one's cognitive character (e.g., Greco 2003; 2010; Sosa 2007; 2011; 2015). Accordingly, cognitive achievements, a special case of achievements more generally, come apart from knowledge in cases of environmental epistemic luck, since they—unlike knowledge—can be unsafe. But one can also have cases of knowledge that don't involve cognitive achievement. The kind of testimonial case that we looked at earlier, where someone is for the most part merely trusting the word of a reliable informant is a case in point. Take, for example, a notable version of such a case, introduced by Jennifer Lackey (2009), where one—having just arrived in a new city—asks a passerby for directions to a major landmark and trusts what this individual says. For while one can plausibly gain knowledge in this fashion, one would hardly suppose that one's cognitive success is *primarily* attributable to one's *own* cognitive ability (as opposed to the informant's cognitive ability), and this is so even if one's exercise of cognitive abilities (e.g., sensitivity to whether the informant looks confused, drunk, etc.) plays some (albeit, not primary³⁵) role in the agent's acquiring the belief that she does.³⁶

The alert reader will see where this is all heading. If understanding is a kind of cognitive achievement, then it is hardly surprising that it stands in the relationship to knowledge that we have identified, since this is the relationship that cognitive achievements bear to knowledge. Moreover, thinking of understanding as a cognitive achievement also explains its distinctive value. After all, achievements are generally considered to be of special value. All things being equal, one would surely prefer a life rich in achievements than a parallel life consisting of lucky successes.³⁷

Consider now the following proposal about cognitive enhancement:

Cognitive Achievement Account of Cognitive Enhancement (CA_{CE}): cognitive enhancement aims at the facilitation of enhanced levels of cognitive achievement.

CA_{CE} has no difficulty explaining why Lestrade exhibits cognitive enhancement even if his knowledge is undermined by environmental epistemic luck, since he would still be exhibiting understanding, which is a kind of cognitive achievement. Even better, CA_{CE} is better accords with the thinking behind cognitive enhancement, which is an improvement in one's cognitive abilities and their role in cognitive success. What is that if not an enhancement in our levels of cognitive achievement? For sure, cognitive achievement often goes hand-in-hand with

knowledge, which is why we might naturally gloss this thought about the nature of cognitive enhancement along the lines of KA_{CE} . But once we are clear how knowledge and cognitive achievements come apart, then it also becomes clear that CA_{CE} is to be preferred as an account of cognitive enhancement.

Finally, notice that CA_{CE} has the virtue of being an *inclusive* account of cognitive enhancement. In order to see this, suppose for a moment that we had used our reflections on the Lestrade case in §2 to motivate an *understanding account of cognitive enhancement* (i.e., the view that cognitive enhancement is geared towards facilitating understanding rather than knowledge). In so doing, we would have been able to account for why Lastrade exhibits cognitive enhancement even in cases where environmental luck impedes knowledge. But we would also be left with a very restrictive account of cognitive enhancement. After all, not all cognitive achievements involve understanding. An operation on one's eyes—or for that matter, any other faculty that has a bearing on cognition³⁸—might well make them much more cognitively effective, such that one not only exhibits greater levels of cognitive success, but also greater levels of cognitive achievement (as these successes are primarily attributable to one's enhanced cognitive performance). But there need be no suggestion that cognitive enhancement of this kind involves greater levels of understanding.³⁹

4. THE AXIOLOGICAL OBJECTION TO COGNITIVE ENHANCEMENT

Thus far we have not registered any claim about whether or to what extent cognitive enhancement *should* be pursued. Rather, what we've suggested is that cognitive enhancement, as such, is best captured with reference to one kind of valuable epistemic standing, cognitive achievement, than it is with reference to another, knowledge. In this section we want to consider an objection, raised in various forms by Leon Kass (2004), Michael Sandel (2009) and John Harris (2011), which threatens to pose a special kind of problem for the cognitive achievement account of cognitive enhancement articulated in the previous section.

What has troubled Kass, Sandel, and Harris with some of the more radical calls for non-traditional cognitive enhancement is the thought that (in short) there is particular value associated with the kind of achievements that involve the overcoming of obstacles. Accordingly, by aiming to remove such obstacles entirely, some of the radical non-traditional forms of cognitive enhancement threaten to diminish a certain valuable dimension of human life.⁴⁰ Each of these writers has captured this point in a different way.

According to Kass (2004, 5), something important is lost when cognitive enhancement (i.e., in the form of drugs) changes how humans function by divorcing their performance from effort. In particular, he suggests that the kind of ‘easy life’ that could be made possible through biotechnology could render what would be otherwise valuable pursuits and achievements, including cognitive achievements, ‘trivial.’⁴¹ Sandel likewise worries that bioenhancement, in the cognitive or physical domain, threatens to undermine the value of human achievement, by disconnecting the success in question from human agency. Consider, for instance, Sandel’s (2012) remarks about bionic athletes and their pharmacologically or genetically aided accomplishments:

[...] as the role of the enhancement increases, our admiration for the achievement fades. Or rather, our admiration for the achievement shifts from the player to his pharmacist [...] This suggests that our moral response to enhancement is a response to the diminished agency of the person whose achievement is enhanced. The more the athlete relies on drugs or genetic fixes, the less his performance represents his achievement. (Sandel 2012, 25-26)

In a similar vein, leading bioethicist Harris, drawing from Milton’s *Paradise Lost*, expresses his concern that bioenhancement threatens the value attained by succeeding in light of our human and thus fallible methods, specifically, by undermining an important freedom: the ‘freedom to fall.’ Writing in response to Savulescu, Harris claims that “Milton’s insight is the crucial role of personal liberty and autonomy: that sufficiency to stand is worthless, literally morally bankrupt, without freedom to fall”, and goes on to argue that “our freedom to fall is ‘precious.’”⁴² (Harris 2011, 110)

Though they express their points in different ways, Kass, Sandel and Harris are concerned that making our cognitive lives too easy by removing the kinds of barriers to cognitive success that are often difficult to overcome would undermine the value of whatever achievements would ensue. Call this the *axiological objection* to cognitive enhancement. A somewhat simplified way of making this point can be made with reference to a computer simulation: suppose you are playing a strategy game that is normally difficult to beat. However, thanks to a ‘cheat code’, all the obstacles from the game are removed—and it is now almost impossible *not* to win. Winning the strategy game, in such a setting, is obviously a kind of success, and perhaps even a kind of achievement too,⁴³ albeit a very minimal one. It doesn’t seem like a particularly valuable success, however. Indeed, it is hard to see why anyone would care about winning such a game.

Another way of appreciating this point will be to focus specifically cases of achievements with what Gwen Bradford (2013) calls ‘zero-value’ products. Consider, for instance, her discussion of the value of summiting a mounting:

[...] achievements that have inherently valueless products gain value in virtue of their difficulty. In mountain climbing, the state of affairs that one aims at—being on top of a mountain—is valuable *only* in the instance when it is attained in a difficult manner, keeping the other features the same. If we get to the top of the mountain by taking a helicopter or an escalator, it's not especially valuable. So it seems that at least in these cases—achievements with zero-value products—difficulty is *responsible* for the value of the achievement. (Bradford 2013, 210, *our italics*)

Interestingly, this way of thinking about the value of achievements, according to which difficulty plays an important role in accounting for their worth, is reflected in the wider philosophical literature, where achievements are often distinguished in terms of the extent to which they involve either high levels of skill or the overcoming of a significant obstacle to success.⁴⁴

On one way of cutting up the landscape which is favoured by Bradford (2013; 2015), we might be inclined to regard difficulty as a *necessary* condition for achievements (as a kind), and not merely as a feature which is a difference maker for the *value* of achievements. While we're sympathetic to this line on which difficulty is an essential feature of achievements, we also think—as do Greco (2010) and Sosa (2009)—that there is something distinctive about non-difficult cases of success through ability that distinguishes these non-difficult cases from their 'Gettiered' success counterparts.⁴⁵

Our terminological preference, which preserves the crux of Bradford's insight that difficulty is an important element of valuable achievements, and which has been defended by one of the authors in previous work (e.g., Pritchard 2009*a*; 2009*b*; 2009*c*; Haddock, Millar & Pritchard 2010, ch. 4), is to distinguish two varieties of achievement, strong and weak. A *strong achievement* is a success due to ability which is difficult, in the sense that it *either* demands high levels of skill or the overcoming of a significant obstacle to success. By contrast a *weak achievement* demands neither, high levels of skill nor the overcoming of significant obstacles to success, but requires only that the success in question be because of ability.⁴⁶ Raising one's arm in normal circumstances—while undoubtedly a success that is primarily attributable to one's relevant abilities—is merely a weak achievement on this taxonomy, since neither of these additional conditions obtain. Raising one's arm when it is being held down by someone, however—perhaps because this person doesn't want you to register your vote—would count as a strong achievement, due to the obstacle to success. Moreover, although the highly skilled might be able to succeed in their fields with ease—consider the sublime shot of a top-ranking tennis player, for example—this still counts as a strong achievement because of the high level of skill involved.

Crucially, strong achievements—and this is a point where we align with Bradford's thinking—are considered more valuable than weak achievements, where the two axes along

which this greater value is measured is in terms of the skill involved and the obstacles that have been overcome. So far, then, this fits very neatly with the objection under consideration, in that if cognitive enhancement leads to a situation in which subjects tend to manifest, at best, only weak cognitive achievements, then this will have deleterious axiological implications. The crux of the matter, however, is whether cognitive enhancement does tend to undermine strong cognitive achievements in this way, as the axiological objection suggests. We want to briefly canvass three reasons why we should be sceptical of this claim, before offering a more detailed exploration of a fourth ground for scepticism in the next section.

The first is that often cognitive enhancement doesn't undermine strong achievements at all, and hence is often very unlike the video-game cheat example that we offered above. Consider, for example, the case of Moddy, who uses Modafinil to increase her focus. This is a eugeroic (i.e., wakefulness-promoting) drug that is ordinarily proscribed for narcolepsy. While under the influence of the drug, Moddy is able to stay awake and concentrate a bit longer than she would be able to without Modafinil, and is thereby able to work out the solution to a math problem on which she would ordinarily have given up (out of exhaustion or lack of focus).⁴⁷

The case of Moddy is analogous to the case of the video-game cheat code only in a very limited respect—*viz.*, that the chance of the relevant success in the two cases, conditioned upon the non-traditional enhancement, is greater than it would be otherwise. But beyond this, there's little similarity. Consider this point first in terms of Moddy's skill. While the Modafinil was clearly an enabling condition for Moddy's success, Moddy's success is creditable to a significant extent nonetheless on her mathematical abilities, without which the problem wouldn't have been solved. It would be a misunderstanding to deny that skill was involved at all in Moddy's success simply because Moddy's naturally endowed skill alone was insufficient without the additional aid in focus and wakefulness.⁴⁸

Second, notice that the use of Modafinil in the case of Moddy doesn't 'remove all obstacles', or for that matter even *most* obstacles.⁴⁹ The mathematical problem is, by stipulation, a complex one. An unskilled mathematician, using Modafinil, would be unable to make much headway. So the presence of cognitive enhancement doesn't remove the scope for exhibiting a strong cognitive achievement.

The third reason for scepticism about the axiological objection is that often cognitive enhancement leads not to weak cognitive achievements but rather to more sophisticated kinds of strong cognitive achievement. For example, if one finds doing complex mathematical puzzles easier due to cognitive enhancement, then one is likely to tackle harder puzzles of this kind. In this way, one is not prompted by the cognitive enhancement to trade strong cognitive

achievements for weak cognitive achievements, but rather to tackle a harder set of strong cognitive achievements. If anything, then, one is led to exhibit cognitive states which are more valuable rather than less.

5. COGNITIVE ENHANCEMENT AND EXTENDED COGNITION

In order to appreciate our fourth basis for scepticism about the axiological objection we first need to make explicit a presupposition which seems to be guiding this thesis. This is that cognitive achievements are to be understood exclusively in terms of one's cognitive skills that are internal to one's biological self—i.e., those cognitive skills of the subject which are rooted in the innate biological traits of the subject, such as her perceptual or rational faculties (call these the *on-board cognitive skills* of the agent). The thinking seems to be that where those faculties are enhanced by external devices, such as wearable technology, then we should only evaluate the subject's cognitive achievements in terms of the contribution of her on-board cognitive skills. Hence, where the cognitive enhancement plays a significant explanatory role in the subject's cognitive success, it immediately follows that the subject is not in the market for a strong cognitive achievement.

A good example of the kind of (non-traditional) cognitive enhancement that is in mind here is what Nick Bostrom and Anders Sandberg (2006, 5) call *intelligence augmentation*.⁵⁰ This includes 'embedding the human within an augmenting "shell" such as wearable computers or virtual reality' and *smart environments*, in which one is surrounded by 'an "exoself" consisting of their files, software, webpages, bookmarks, online identities and other personal information.' One very natural example of non-traditional cognitive enhancement *via* intelligence augmentation is Google Glass 'smartglasses' with feature an optical head mounted display allowing for 'ubiquitous computing.'⁵¹

To use a concrete example, consider a recent case study published by Oliver Muensterer *et al* (2014) in which doctors at a children's hospital used Google Glass for four consecutive weeks. In this study a daily log was kept, and activities with a potential applicability were identified, and performance of Google Glass was evaluated for such activities. Participants in the study reported that, among other things, they worked more effectively and quickly with information related to (previously) unfamiliar medical terms or syndromes. A related and even more striking example is Phillips Healthcare's new use of Google Glass to overlay information

into the clinician's field of view, in such a way as to allow doctors to continuously monitor patients' vital signs during surgical procedures without ever taking their eyes off the patient.⁵²

Let's simply grant that intelligence augmentation via Google Glass can afford individuals capabilities which outstrip their naturally (i.e., biologically) endowed abilities. We can restate the reformulated axiological objection so as to target cases like these—by claiming that in a wide class of 'intelligence augmented cognitive successes', the success is credible to the technology more so than to the agent, in such a way that what is left is not something to which we can properly credit the agent. If this line of thinking is right, then a wide class of enhanced cognitive successes are outside the class of valuable *human* achievements in that they are not significantly attributable to the agents concerned (unlike strong cognitive achievements).

Interestingly, however, recent work in the philosophy of mind and cognitive science offers some resources for undermining this line of argument. In short, the matter of whether a given intelligence augmentation device (and, more generally, a non-biological artifact) should be regarded as part of one's own agency, or external to it, depends crucially on whether that item has been appropriately integrated within the individual's cognitive system. So, for example, proponents of *extended cognition*—*viz.*, the thesis that some cognitive processes can crisscross the boundaries of brain and world—have claimed that technology which is appropriately embedded within a subject's cognitive system, such that there are systematic feedback loops at issue between the subject's onboard cognition and the embedded technology, should be regarded as a genuine part of the subject's cognitive system.⁵³ Moreover epistemologists have argued that in cases of extended cognition, the cognitive success in play should be treated as attributable to the cognitive subject, and hence that the cognitive subject should be ascribed the relevant epistemic states (e.g., cognitive achievement, knowledge, understanding, and so on).⁵⁴

It would obviously take us too far afield to defend either the extended cognition programme or its putative epistemological ramifications here. Our key point, however, is that proponents of the axiological objection are effectively presupposing that cognitive enhancement will not go down this route, and yet that is, in fact, a highly controversial supposition in this context. This is yet another reason for being sceptical about the import of this objection.

6. CONCLUDING REMARKS

We conclude that the proper way to think about the nature of cognitive enhancement is in terms of the facilitation of cognitive achievements rather than knowledge. Moreover, we contend that

concerns about the desirability of such enhancement, at least as currently articulated, are far from convincing. In particular, there are several grounds available for resisting the charge that cognitive enhancement is bound to have negative axiological consequences.⁵⁵

REFERENCES

- Alston, W. P. 2005. *Beyond justification: dimensions of epistemic evaluation*, Ithaca, NJ: Cornell University Press.
- Battaly, H. 2008. Virtue epistemology. *Philosophy compass* 3: 639–63.
- Beebe, J. 2004. Reliabilism, truetemp and new perceptual faculties. *Synthese* 140: 307–29.
- Bonjour, L. 1980. Externalist theories of empirical knowledge. *Midwest Studies in Philosophy* 5: 53–73.
- Boorse, C. 1975. On the distinction between disease and illness. *Philosophy and Public Affairs* 5: 49–68.
- Bostrom, N. & Roache, R. 2007. Human enhancement: ethical issues in human enhancement. In: *New Waves in Applied Ethics*, J. Ryberg, T. Petersen, & C. Wolf (eds). London: Palgrave Macmillan.
- Bostrom, N. & Savulescu, J. 2009. Human enhancement ethics: the state of the debate. In: *Human Enhancement* (pp. 1–22), N. Bostrom & J. Savulescu (eds). Oxford: Oxford University Press.
- Bradford, G. 2013. The value of achievements. *Pacific philosophical quarterly* 94: 204–224.
- 2015. *Achievement*. Oxford: Oxford University Press.
- Bradford, G. & Keller, S. *Forthcoming*. Achievement and well-being. In: *Routledge Handbook of Philosophy of Well-Being*, G. Fletcher (ed.). London: Routledge.
- Buchanan, A. 2008. Enhancement and the ethics of development. *Kennedy Institute of Ethics Journal* 18: 1–34.
- Carter, J. A. & Gordon, E. C. 2013. Intelligence, wellbeing and procreative beneficence. *Journal of Applied Philosophy* 30: 122–35.
- 2014. On cognitive and moral enhancement: a reply to Savulescu and Persson. *Bioethics* 29: 153–161
- Carter, J. A., Jarvis, B. W., & Rubin, K. 2015. Varieties of cognitive achievement. *Philosophical Studies* 172: 1603–1623.
- Carter, J. A. & Pritchard, D. H. 2015a. Knowledge-how and epistemic luck. *Noûs*, 49: 440–453.
- 2015b. Perceptual knowledge and relevant alternatives. *Philosophical studies* 173: 969–990.
- Carter, J. A., Kallestrup, J., Palermos, S. O., & Pritchard, D. H. 2014. Varieties of externalism. *Philosophical Issues*, 24: 63–109.
- Clark, A. 2008. *Supersizing the mind: embodiment, action, and cognitive extension*. Oxford: Oxford University Press.
- Clark, A., Carter, J. A., Kallestrup J., Palermos, S. O., & Pritchard, D. H. 2014. Varieties of externalism. *Philosophical Issues* 24: 63–109.
- Clark, A. & Chalmers, D. 1998. The extended mind. *Analysis* 58: 7–19.
- Craig, E. 1990. *Knowledge and the state of nature*. Oxford: Oxford University Press.
- Doris, J. 1998. Persons, situations, and virtue ethics. *Noûs* 32: 504–30.
- 2002. *Lack of character: personality and moral behavior*. Cambridge: Cambridge University Press.
- Egan, G. 1995. *Permutation City*. New York: Eos.
- Elgin, C. 2009. Is understanding factive?. In: *Epistemic value*. A. Haddock, A. Millar & D. H. Pritchard (eds.). Oxford: Oxford University Press.
- Feinberg, J. 1970. *Doing & deserving: essays in the theory of responsibility*. Princeton, NJ: Princeton University Press.
- Fitz, N. S., Nadler, R., Manogaran, P., Chong, E. W. & Reiner, P. B. 2014. Public attitudes toward cognitive enhancement. *Neuroethics* 7: 173–188.
- Hurka, T. 1993. *Perfectionism*. Oxford: Oxford University Press.
- Glüer, K. & Wikforss, Å. 2009. Against content normativity. *Mind* 118: 31–70.

- Goldman, A. 2011. Reliabilism. In: *Stanford Encyclopedia of Philosophy*, (ed.) E. Zalta, <http://plato.stanford.edu/archives/spr2011/entries/reliabilism/>.
- Goldman, A. & Olsson, E. J. 2009. Reliabilism and the value of knowledge. In: *Epistemic Value* (pp. 19–41), A. Haddock, A. Millar & D. H. Pritchard (eds), 19–41. Oxford: Oxford University Press.
- Greco, J. 2010. *Achieving knowledge*. Cambridge: Cambridge University Press.
- 2012. A (different) virtue epistemology. *Philosophy and phenomenological research* 85: 1–26.
- Greco, J. & Turri, J. 2011. Virtue epistemology. In: *Stanford Encyclopedia of Philosophy*.
- Grimm, S. 2010. Understanding. In: *Routledge Companion to Epistemology*. S. Bernecker & D. H. Pritchard (eds). New York: Routledge.
- Harris, J. 2011. Moral enhancement and freedom. *Bioethics* 25: 102–11.
- Hoffman, P. 1998. *The man who loved only numbers: the story of Paul Erdős and the search for mathematical truth*. London: Fourth Estate.
- Hauskeller, M. 2013. Cognitive enhancement—to what end?. In: *Cognitive enhancement*. E. Hildt & A. G. Franke (eds.) (pp. 113–23). Netherlands: Springer.
- Juengst, E. & Moseley, D. 2015. Human enhancement. In: *Stanford Encyclopedia of Philosophy*. E. Zalta (ed). <http://plato.stanford.edu/archives/sum2015/entries/enhancement/>.
- Kallestrup, J., & Pritchard, D. H. 2014. Virtue epistemology and epistemic twin earth. *European journal of philosophy* 22: 335–357.
- Kass, L. R. 2004. *Life, liberty and the defense of dignity: the challenge for bioethics*. San Francisco, CA: Encounter Books.
- Kelp, C. Forthcoming. Understanding phenomena. *Synthese*.
- 2014. Two for the knowledge goal of inquiry. *American philosophical quarterly* 51: 227–232.
- Kitcher, P. 2002. Scientific knowledge. *Oxford handbook of epistemology*. P. Moser (ed.) Oxford: Oxford University Press.
- Kvanvig, J. 2003a. *The value of knowledge and the pursuit of understanding*. Cambridge: Cambridge University Press.
- 2003b. Simple reliabilism and agent reliabilism. *Philosophy and phenomenological research* 66: 451–56.
- 2009. The value of understanding. In: *Epistemic Value* (pp. 95–112), A. Haddock, A. Millar, & D. H. Pritchard (eds.). Oxford: Oxford University Press.
- Karttunen, L. 1977. Syntax and semantics of questions. *Linguistics and philosophy* 1: 3–44.
- Lackey, J. 2009. Knowledge and credit. *Philosophical studies* 142: 27–42.
- Lehrer, K. 2000. *Theory of knowledge* (2nd Edn.). London: Routledge.
- Lipton, P. 2004. *Inference to the best explanation*. New York: Routledge.
- Mankar, P. 2015. Advance technology: google glass. *International research journal of engineering and technology* 2: 73–77.
- Montmarquet, J. 1993. *Epistemic virtue and doxastic responsibility*. London: Rowman & Littlefield.
- Muensterer, O. J., Lacher, M., Zoeller, C., Bronstein, M. & Kübler, J. 2014. Google glass in pediatric surgery: an exploratory study. *International Journal of Surgery* 12: 281–89.
- Müller, U., Rowe, J. B., Rittman, T., Lewis, C., Robbins, T. W. & Sahakian, B. J. 2013. Effects of modafinil on non-verbal cognition, task enjoyment and creative thinking in healthy volunteers. *Neuropharmacology* 64: 490–495.
- Palermos, S. O. 2011. Belief-forming processes, extended. *Review of philosophy and psychology* 2: 741–65.
- 2014. Loops, constitution, and cognitive extension. *Cognitive systems research* 27: 25–41.
- Persson, I. & Savulescu, J. 2008. The perils of cognitive enhancement and the urgent imperative to enhance the moral character of humanity. *Journal of applied philosophy* 25: 162–77.
- 2012. *Unfit for the future: the need for moral enhancement*. Oxford University Press.
- Plomin, R. 1997. Identifying genes for cognitive abilities and disabilities. In: *Intelligence, Heredity, and Environment* (pp. 89–104), R. J. Sternberg & E. L. Grigorenko (eds.). New York:

- Cambridge University Press.
- Pritchard, D. H. 2007. Recent work on epistemic value. *American philosophical quarterly* 44: 85-110.
- 2009a. Apt performance and epistemic value. *Philosophical Studies* 143: 407-16.
- 2009b. *Knowledge*. London: Palgrave Macmillan.
- 2009c. Knowledge, understanding and epistemic value. In: *Epistemology (Royal Institute of Philosophy Lectures)*, (pp. 19–43), A. O’Hear (ed.). Cambridge: Cambridge University Press.
- 2010a. Achievements, luck and value. *Think* 25: 1-12.
- 2010b. Cognitive ability and the extended cognition thesis. *Synthese* 175: 133-51.
- 2012. Anti-luck virtue epistemology. *Journal of philosophy* 109: 247-79.
- 2014. Knowledge and understanding. In: *Virtue Scientia: Bridges Between Virtue Epistemology and Philosophy of Science* (pp. 315-28), A. Fairweather (ed.). Dordrecht, Holland: Springer.
- 2015a. Anti-luck epistemology and the gettier problem. *Philosophical studies* 172: 93-111.
- 2015b. Risk. *Metaphilosophy* 46: 436-61.
- Forthcoming. Epistemic risk. *Journal of Philosophy*.
- 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. In: *Stanford encyclopædia of philosophy*, E. Zalta (ed.). <http://plato.stanford.edu/entries/knowledge-value/>.
- Riggs, W. 2007. Understanding ‘virtue’ and the virtue of understanding. In: *Intellectual virtue: perspectives from ethics and epistemology*, M. DePaul & L. Zagzebski (eds.). Oxford: Clarendon Press.
- Sandberg, A. & Bostrom, N. 2006. Cognitive enhancement: a review of technology. <http://diyhl.us/~bryan/papers2/neuro/implants/Anders%20Sandberg,%20Nick%20Bostrom%20-%20Cognitive%20Enhancement%20Tech%20Review.pdf>.
- Sandel, M. 2012. The case against perfection: what’s wrong with designer children, bionic athletes, and genetic engineering?. In: *Arguing About Bioethics*, S. Holland (ed.). London: Routledge.
- Savulescu, J. 2001. Procreative beneficence: why we should select the best children. *Bioethics* 15: 413–26.
- Savulescu, J., ter Meulen, R. & Kahane, G. (eds.) 2011. *Enhancing human capacities*. New York: Wiley-Blackwell.
- Savulescu, J. & Person, I. 2012. Moral enhancement, freedom and the god machine. *The monist* 95: 399–421.
- Shah, N. & Velleman, J. D. 2000. Doxastic deliberation. *Philosophical review* 114: 497–534.
- Sliwa, P. 2015. Understanding and knowing. *Proceedings of the Aristotelian Society* 115: 57-74.
- Sosa, E. 2000. For the love of truth. In: *Virtue epistemology: essays on epistemic virtue and responsibility* (pp. 49-62), L. Zagzebski (ed.). Oxford: Oxford University Press.
- 2007. *Reflective knowledge: apt belief and reflective knowledge*. Oxford: Oxford University Press.
- 2011. *Knowing full well*. Princeton, NJ: Princeton University Press.
- 2009. *A virtue epistemology: apt belief and reflective knowledge*. Oxford: Oxford University Press.
- 2015. *Judgment and agency*. Oxford: Oxford University Press.
- Stanley, J. 2011. *Know how*. Oxford: Oxford University Press.
- Strevens, M. 2013. No understanding without explanation. *Studies in History and Philosophy of Science* 44: 510–15.
- Terbeck, S., Kahane, G., McTavish, S., Savulescu, J., Cowen, P. & Hewstone, M. 2012. Propranolol reduces implicit negative racial bias. *Psychopharmacology* 222: 419–24.

- Thagard, P. 2014. Cognitive science. In: *Stanford encyclopedia of philosophy*. E. Zalta (ed.). <http://plato.stanford.edu/archives/fall2014/entries/cognitive-science/>.
- Thomas, J. & Brunnsma, D. 2014. Oh, you're racist? I've got a cure for that!. *Ethnic and racial studies* 37: 1467–85.
- Turri, J. 2015. Epistemic situationism and cognitive ability. In: *Epistemic situationism*, M. Alfano & A. Fairweather (eds.). Oxford: Oxford University Press.
- Forthcoming. Knowledge as achievement, more or less. *Performance epistemology*, M. A. Fernandez (ed.). Oxford: Oxford University Press.
- 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 For a helpful contemporary overview of some of contemporary arguments for and against various kinds of human enhancement, see Juengst & Moseley (2015). See also Savulescu *et al* (2011) for some influential recent essays.

² In the latter case, the most substantial argument on this score can be found in Persson & Savulescu (2008; 2012). See Harris (2011) and Carter & Gordon (2015) for criticisms.

³ The cognitive is often contrasted with the *affective* (i.e., emotion) and the *conative* (i.e., volition), as separate aspects of the mind. The scientific study of cognition, cognitive science, is multi-disciplinary (see, e.g., Thagard 2014), and within cognitive science and the philosophy of mind, the question of whether there is a metaphysical ‘mark of the cognitive’ is itself controversial and at any rate not a point we will be taking a stand on here.

⁴ For further discussion of this distinction, see Bostrom & Sandberg (2009, 312).

⁵ See Boorse (1975) for a helpful discussion of the difference between disease and health. Cf., Sandel (2012, 8).

⁶ For a recent overview of public attitudes towards non-traditional forms of cognitive enhancement, see Fitz *et al* (2014).

⁷ One live debate in this regard centres around how to conceive of normalcy. For instance, is the notion of normalcy relevant to distinguishing between therapeutic and non-therapeutic enhancement normalcy relative to what is typical of the individual in question, or rather, normalcy relative to what is statistically within a certain average for the individual’s (relevant) population?

⁸ For example, cochlear implants. Neuroprosthetics differ from brain-computer interfaces, in that neuroprosthetics are meant to replace or enhance brain functionality whereas brain-computer interfaces put the brain in contact with an external computer.

⁹ See Bostrom & Savulescu (2009) and Persson & Savulescu (2012).

¹⁰ See Bostrom & Roache (2007, 226).

¹¹ As Bostrom & Savulescu (2009, 2) ask:

“How is taking modafinil fundamentally different from imbibing a good cup of tea? How is either morally different from getting a full night’s sleep? Are not shoes a kind of foot enhancement, clothes an enhancement of our skin? A notepad, similarly, can be viewed as a memory enhancement—it being far from obvious how the fact that a phone number is stored in our pocket instead of our brain is supposed to matter once we abstract from contingent factors such as cost and convenience.”

¹² As they put it, ‘the most likely effect of cognitive enhancement by genetic or biomedical means may be to speed up a growth of *knowledge* that would otherwise have taken humanity a longer time to achieve.’ (Persson & Savulescu 2008, 7)

¹³ There are several ways one might take such a line. One might be a nihilist about cognitive norms, for reasons broadly akin to why some philosophers (e.g., Glüer-Pagin & Wikforss 2009) deny that belief and content are in any way normative. Alternatively, one might think that cognitive value is exhausted by *responsible* inquiry, and independent of the kinds of representational states that are attained. One branch of contemporary virtue epistemology that places a premium on responsibility is *virtue responsibilism* (e.g., Montmarquet 1993; Battaly 2008). However, virtue responsibilists are free to embrace different kinds of accounts (e.g., personal worth accounts, such as Baehr 2011, or teleological accounts, such as Zagzebski 2006) about what makes responsible inquiry a cognitive good. See for example Greco & Turri (2011) for an overview.

¹⁴ In epistemology, one popular position is that the aim of cognition is *truth*, a view that has been embraced under the description of epistemic value truth-monism. See, for example, Lynch (2009) and Sosa (2000) for defences of this position. A rival view, one that has been embraced by Craig (1990), Henderson (2009; 2011), Kelp (2014) and by proponents of the ‘knowledge-first’ approach championed by Williamson (2000), insists that knowledge is the governing value around which cognitive assessments should be structured. For discussion (more generally) of how the identification of a primary cognitive value or aim structures the normativity of cognitive appraisals, see Alston (2006).

¹⁵ For an overview of literature on this topic, see Kvanvig (2003), Pritchard & Turri (2014), Haddock, Millar & Pritchard (2010).

¹⁶ This is a variation on the famous ‘TrueTemp’ case offered by Lehrer (2000). For some other biological variants on this case, see Beebe (2004).

¹⁷ According to Bonjour (1980), the protagonist in such cases is ‘subjectively irrational’ in holding the target belief. See also Goldman (2011, §3).

¹⁸ Suppose that *Chip**’s chip is thoroughly integrated into *Chip**’s parietal lobe, as well as more *generally* into *Chip**’s cognitive system—in such a way that *Chip**’s directional orientation beliefs are not *merely* reliable (as in the case of *Chip*), but also that the resultant true beliefs are beliefs for which *Chip* can provide rational grounds.

¹⁹ See, for example, Bonjour (1980), Lehrer (2000), Pritchard (2010*b*), and Goldman (2011).

²⁰ The proponent of TB_{CE} might attempt to make a rebuttal here to the effect that the right explanation for why *Chip** is enhanced to a greater extent than *Chip* is that the former is better equipped to attain more true beliefs *in the future* than the latter is, owing to his greater degree of rationality or coherence. The explanation, however, merely forestalls the problem for TB_{CE}. Consider that even if, *ex hypothesi*, *Chip* were to acquire the very same true beliefs

which Chip* acquires in the future, Chip* is in a better epistemic position. But, if TB_{CE} were true, they would on this assumption be cognitively enhanced to an equal extent, which is false. For a related discussion of epistemic value in terms of the production of future true beliefs, see Goldman & Olsson (2009). Thanks to an anonymous referee for the *Journal of Medicine and Philosophy* for raising this point.

²¹ Savulescu (2001, 413, 421) argues that the potential for genetic forms of cognitive enhancement is rapidly progressing in virtue of the development of what is known as *molecular genetic studies*, the aim of which is to locate and characterise specific genes that play a biological role in intelligence. See Plomin (1997) for some criticisms. See also Carter & Gordon (2013) for an overview.

²² For more on the nature of understanding, and the various different kinds of understanding that can be distinguished, see Grimm (2010).

²³ See, for example, Kitcher (2002), Lipton (2004), and Strevens (2013).

²⁴ See Sliwa (2015) and Kelp (*forthcoming*) for two recent exceptions.

²⁵ See, for example, Pritchard (2015*a*) for a recent articulation and defence of this claim.

²⁶ For further discussion of the notion of epistemic risk in this regard, and its relationship to the notion of risk more generally, see Pritchard (2015*b*; *forthcoming*).

²⁷ This distinction is developed and explored in Pritchard (2009*a*, 2009*b*, 2009*c*; 2012) and Haddock, Millar & Pritchard (2010, chs. 1-4). See also Pritchard & Kallestrup (2014), which re-articulates the distinction in terms of an epistemic twin earth argument.

²⁸ These cases originally appeared in Pritchard (2009*b* ch. 7). See also Pritchard (2009*c*; 2014) and Haddock, Millar & Pritchard (2010, ch. 4).

²⁹ For some examples of authors who argue for the distinctive value of understanding, see Kvanvig (2003; 2009), Riggs (2007), Grimm (2010), Pritchard (2009*c*; 2014), and Haddock, Millar & Pritchard (2010, chs. 1-4). For two recent surveys of work on the nature of epistemic value and its import to contemporary epistemology, see Pritchard (2007) and Pritchard & Turri (2011).

³⁰ Might the proponent of KA_{CE} attempt to sidestep this objection by pointing to Lestrade's capacity to acquire future instances of knowledge, even if not acquiring knowledge on the present occasion? Such an explanation is committed to the result that Lestrade, were his enhancement to facilitate the acquisition of future knowledge which is the same future knowledge facilitated by an enhancement to his counterpart who also attains understanding, that Lestrade and his understanding attaining counterpart would be equally enhanced. We deny this. Of course, our denial would count for naught if sameness of knowledge entailed sameness in understanding. Though, we've argued in the foregoing that it does not.

³¹ Thanks to an anonymous referee for the *Journal of Medicine and Philosophy* for requesting clarification on this point.

³² For a recent articulation, and defence, of the cognitive achievement account of understanding, see Pritchard (2014).

³³ Greco's account of cognitive achievement gives the notion of 'credit' a central place, such that achievements as such are creditable successes. Greco's discussion of intellectual credit is informed by Feinberg's (1970) account of moral credit attributions.

³⁴ The precedent in the virtue epistemology literature, which we follow here, is to use 'true belief' and 'cognitive success' interchangeably. The reason for stressing 'cognitive success' is to more clearly highlight the role true belief is playing in the structure of cognitive achievements, which are special cases of achievements involving success and ability, more generally.

³⁵ For discussion on this point, see Pritchard (2012). See Greco (2012) for a reply to such criticisms. See, however, Carter, Jarvis & Rubin (2015) for a counterreply to Greco regarding testimony-style objections to the cognitive achievement account of propositional knowledge embraced by Greco and other robust virtue epistemologists.

³⁶ And of course what goes for knowledge here also applies, *a fortiori*, to mere true belief. For example, if one found a way of simply implanting true beliefs into a subject, whereby the cognitive success in question bore absolutely no explanatory relationship to that subject's manifestation of relevant cognitive agency, then this wouldn't count as a cognitive achievement on the part of the subject.

³⁷ For more discussion of the value of achievements in general, see Pritchard (2010*a*).

³⁸ While we are contrasting the cognitive from the effective and conative, we should note that, in principle at least, various kinds of human faculties could potentially be responsible for human cognizing. While the eyes directly contribute to our perceptual belief formation, so do other senses contribute to (for instance) auditory and olfactory belief formation. An interesting question concerns to what extent the improvement of non-biological elements of the world we interact with might also constitute cognitive enhancement, and it's a question which depends on our accepted theory of the bounds of cognizing and whether these bounds are biologically constrained. See our discussion in §5 on extended cognition, which takes up some of these issues.

³⁹ Note that on this view the implantation of true beliefs into a subject would not in itself guarantee that the subject has been cognitively enhanced, since this would further depend on whether those true beliefs facilitate enhanced levels of cognitive achievement.

⁴⁰ See Hurka (1993) for a statement of the view that achievement is an element of a valuable life within the purview of a perfectionist model of human well-being.

⁴¹ For further discussion, see the 2007 white paper from the British Medical Association, ‘Boosting your Brainpower: Ethical Aspects of Cognitive Enhancements. A Discussion Paper from the British Medical Association’ (<http://hdl.handle.net/10822/511709>).

⁴² See Persson & Savulescu (2012) for a reply to Harris’s deployment of this argument in the service of criticising specifically Persson & Savulescu’s moral argument for non-traditional methods of moral bioenhancement.

⁴³ This will not be the case, however, on models such as Bradford’s (2013) according to which difficulty is an essential element of achievement, as such which are characterized by ‘a difficult process which culminates competently in a product’ (2013, 205). On the model we are sketching, strong cognitive achievements have a difficulty requirement while weak cognitive achievements do not.

⁴⁴ See, for example, Bradford (2013; 2015), Bradford & Keller (*forthcoming*), Greco (2010), and Feinberg (1970).

⁴⁵ ‘Gettiered’ in the sense of intervening luck. We regard achievements as compatible with environmental luck.

⁴⁶ See Pritchard (2010*a*) for further discussion of this distinction in general, and Haddock, Millar & Pritchard (2010, ch. 4) for further discussion of this distinction in the specific context of cognitive achievements.

⁴⁷ A recent study by Müller *et al* (2013) suggests that Modafinil also can increase self-reported enjoyment in non-verbal cognitive tasks among healthy individuals. Incidentally, a real-life example in the neighbourhood of this case involves the Hungarian mathematician Paul Erdős who, according to his biographer Paul Hoffman (1998, ch. 1), ‘put in nineteen-hour days, keeping himself fortified with 10 to 20 milligrams of Benzedrine or Ritalin, strong espresso, and caffeine tablets’.

⁴⁸ Compare this point to the case of Jamaican sprinter Usain Bolt, whose personal best in the 100 metre dash at a time of 9.58 seconds is a world record. It’s not implausible that were Bolt to have run his 9.58 second dash in standard issue shoes from the 1970s rather than with his state of the art Puma Ignite shoes, this could have cost him (at least) .001 of a second, and accordingly he couldn’t have run the 100 metres in 9.58. Granting this point, however, is entirely compatible with attributing Bolt’s success largely to his own ability. (After all, no one else in the world can run 100 metres in 9.58 seconds, with or without Puma Ignite shoes).

⁴⁹ It is tempting to wonder: what if Modafinil were capable of removing *all* obstacles to the mathematical problem in question? In such a circumstance, would the individual’s solving a given problem potentially qualify as a strong cognitive achievement or not? We should note here that it is very plausibly a conceptual truth that if *all* obstacles to a given success are removed, then the success in question cannot be a strong cognitive achievement, as we’ve defined it. The lack of *any* obstacles that might stand in the way of success is after all incompatible with such a success requiring the overcoming of obstacles or for that matter the execution of great skill. However, even if an individual ‘outsourced’ all her naturally endowed capacities for wakefulness and attention to Modafinil, in light of the drug’s function, this would still be compatible with strong cognitive achievement given the extent to which Modafinil is significantly *limited* in which obstacles to the success of a mathematical problem can be removed. Modafinil can potentially remove *focus* based obstacles to solving a mathematical problem, though focus based obstacles are just one of many potential obstacles to solving such a problem. We thank an anonymous referee for the *Journal of Medicine and Philosophy* for raising this question.

⁵⁰ Sandberg & Bostrom borrow this term from Egan (1995).

⁵¹ For a helpful recent overview of Google Glass functionality, see Mankar (2015).

⁵² See <http://www.healthcare.philips.com/main/about/future-of-healthcare>.

⁵³ The *locus classicus* in this regard is Clark & Chalmers (1998). See also Clark (2008). Note that the thesis of extended cognition is in fact weaker than the closely related *extended mind thesis*, which claims that some mental states supervene on extra-organismic elements of the world. This distinction is relevant because one notable objection to the extended mind thesis—known as the *cognitive bloat objection*—cannot very easily be redeployed against the extended cognition thesis. For more on this point, see Carter *et al* (2014). In any case, for our purposes what’s relevant is just that the extended cognition thesis provides resources for mitigating the force of the envisioned axiological objection, and not that the comparatively more radical extended mind thesis provides such resources. For more on the distinction between the extended cognition and extended mind theses, see Palermos (2011; 2014).

⁵⁴ See, for example, Pritchard (2010*b*). For a survey of the issues with regard to the epistemological implications of extended cognition, see Carter *et al* (2014).

⁵⁵ Thanks to Emma C. Gordon and Rie Iizuka for helpful comments on a previous version of this paper. Thanks also to three anonymous referees at the *Journal of Medicine and Philosophy*.