THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Neuromedia and the epistemology of education

Citation for published version:

Pritchard, D 2018, 'Neuromedia and the epistemology of education', *Metaphilosophy*, vol. 49, no. 3, pp. 328-349. <https://doi.org/10.1111/meta.12295>

Digital Object Identifier (DOI):

[10.1111/meta.12295](https://doi.org/10.1111/meta.12295)

Link:

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

Document Version:

Peer reviewed version

Published In:

Metaphilosophy

Publisher Rights Statement:

This is the peer reviewed version of the following article: Pritchard, D. (2018), Neuromedia and the Epistemology of Education. *Metaphilosophy*, 49: 328-349. doi:10.1111/meta.12295, which has been published in final form at <https://onlinelibrary.wiley.com/doi/abs/10.1111/meta.12295>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Self-Archiving.

General rights

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

Take down policy

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

For *Metaphilosophy*

(special issue on 'Connecting Virtues: Theoretical and Educational Insights',
(eds.) M. Croce & M. S. Vaccarezza).

Neuromedia and the Epistemology of Education

Duncan Pritchard

University of California, Irvine & University of Edinburgh

dhpritch@uci.edu & duncan.pritchard@ed.ac.uk

ABSTRACT. I explore the implications of a technological revolution that many in the industry think is likely to soon come to pass: *neuromedia*. In particular, I'm interested in how this will constitute an especially persuasive kind of extended cognition, and thereby will facilitate extended epistemic states. I think this will in turn have ramifications for how we understand the epistemic goals of education. I will be arguing that the challenges posed by neuromedia remind us that the overarching epistemic goal of education is not orientated towards facilitating a body of knowledge (or the development of mere cognitive skills), but is rather concerned with the development of *intellectual character*, where this in turn essentially involves the cultivation of intellectual virtues, character-traits that are not amenable to extended cognition.

KEYWORDS: Education; Epistemology; Extended Cognition; Intellectual Character; Intellectual Virtue.

1. NEUROMEDIA

The technological advances of the last couple of decades have been staggering, and for many people in the world have had profound affects on our everyday lives. Moreover, the pace of technological development is, if anything, gaining momentum all the time. I want to explore one particular kind of technological development that might well be right around the corner:

neuromedia.¹ By neuromedia I have in mind the development of information processing technology that is so seamlessly integrated with our on-board cognitive processes that the subject is often unable to distinguish between her use of those on-board processes and the technology itself. The subject's relationship to the technology is consequently no longer one of subject-to-instrument, but rather 'feels' like a technological extension of her normal cognitive processes.

In order to see how this might work, consider the recent changes in how we access information. Suppose one needs to know the answer to a specific question, such as how many moons Saturn has. At one time, finding this out might have required a trip to the library, or at least phoning up a more knowledgeable friend (at one point it might have even required being able to

design and build a rudimentary telescope). These days, of course, one can just Google this question on one's phone and get the answer within seconds. Technological developments have therefore made it much easier to gain information. Even though looking up an answer on your phone is a far easier way of finding something out than visiting a library, one's relationship to the phone is still one of subject and instrument, however. You are conscious, after all, of using the phone to find this out; that it was the phone that communicated to you the answer to your question.

Imagine, however, that one was technologically augmented in such a way that when a question like this occurs to one, then the answer becomes immediately present to mind (we don't need to worry just now about how this is done). In particular, imagine that the accessibility of the answer, and the associated phenomenology involved, is just like remembering this answer yourself. If technology could be made to work in this fashion, then this would be neuromedia in the sense that I have in mind. In particular, one's relationship to the technology would not now be essentially one of subject-and-instrument (even though one is in fact using an instrument, in effect), in that one might not even be aware that one is employing the technology.

Neuromedia, if it happens, will constitute a transformational shift in our relationship to technology. A wealth of factual knowledge—if knowledge is what it is; we will return to this point—will potentially be at our fingertips, phenomenologically on a par with knowledge that we have personally acquired. Moreover, a range of skills that are currently prized will start to become redundant, as they are off-loaded onto technology. After all, many of our current skills—whether navigational, memorial, arithmetical, linguistic, etc.—are such that they will be more effectively managed via the neuromedia. If this happens, it will have seismic effects on our society.

For the remaining sections of this paper I am going to assume that neuromedia is on the horizon in order to explore its philosophical implications, especially in epistemology, and in particular for the epistemology of education. Before I do so, however, I want to make some brief remarks about this kind of technology, and also flag some broader issues that it raises.

First off, notice that there is more than one way that this kind of cognitive augmentation could happen.² One might imagine that this would be achieved via technological implantment—i.e., the *cyborg* route. The very name 'neuromedia' implies this, since it suggests that one is adding technological media to the cognising subject. One difficulty with this route, however, is that it may be prohibitively expensive, as the augmentation needs to be replicated for every subject. Indeed, one could imagine a dystopian future in which only the super-rich are cognitively augmented in this way, and hence are not only a class apart from everyone else in terms of their wealth but also in terms of their cognitive powers.³

Another way for neuromedia to develop is by embedding subjects within high-tech adaptive environments (though this will often involve some degree of individual cognitive augmentation as well, albeit of a lesser degree to that found with the cyborg route). This is sometimes referred to as the *Minority Report* route because it mirrors how the technology works in that film (which is now famous for predicting lots of new technological developments). The thought is that if one's environment is enriched with technology that is responsive to cognitive subjects, thereby providing each subject a 'bespoke' cognitive environment, then this lessens the need for individual cognitive augmentations. (In the film, for example, the environment would recognise the subject when present and offer information that it thought would be relevant for that particular person). This option may prove cheaper overall, but it also has a lot of downsides (some of which are depicted in *Minority Report* itself). After all, it is unlikely to be the state that creates this structural environmental technology for its citizens, in which case it will be in the hands of corporations who may not have individual citizens best interests at heart. (In the film this technology essentially gave people individually tailored adverts, of the kind one is now familiar with on, say, Facebook).

This last point reminds us that even if we can overcome the technological hurdles facing the development of neuromedia, there may be other barriers to its implementation. For example, there might be legal concerns. If subjects are unable to know when they are relying on their on-board cognitive resources rather than the technology, then this is likely to have important implications when it comes to issues like legal liability. For example, if it is the technology that is guiding one's calculations, and this has a bug that leads to a mistake with important legal consequences, then in what sense are *you* liable for this error (why not the company that installed the technology)? Or, to take another example, if I testify to something under oath, and it turns out that the source of this information is (unbeknownst to me) the technology rather than my biological memory, then have I misled the court? Relatedly, should technology-assisted testimony even count as admissible evidence in court?

There are also obvious political concerns regarding the misuse and regulation of this kind of technology. It can clearly be used to deceive us, and thereby deny us epistemic goods, as much as it can be used to cognitively enhance us. And if cognitive enhancement is available, who is to receive it? We've already noted the dangers of a specifically cognitive inequality, whereby there emerges a cognitive underclass. But is it even feasible to provide such augmentations to everyone? Could one have a *right* to be cognitively augmented (in the way that many hold that one has a right to access to a good education, to good health services, and so on)? And what about those who don't want to be cognitively augmented? Could we imagine a society that obliges everyone to be cognitively augmented, so that no-one is left behind?

These are all fascinating questions, but we will be setting them to one side in what follows. Our focus will instead be on understanding the epistemological ramifications of neuromedia, especially with regard to the epistemology of education. In particular, we will be assuming that neuromedia will constitute a genuine cognitive augmentation of the subject (i.e., and not be used to manipulate or deceive us, etc.).

2. EXTENDED COGNITION

Neuromedia is a form of *extended cognition*.⁴ Indeed, I submit that it is the *most* plausible candidate for this title. According to extended cognition, a subject's cognitive processes can extend beyond her brain and central nervous system; indeed, can extend beyond her skin and skull. So, for example, features of the subject's environment could be employed in such a way that they become genuine parts of the cognitive process itself (i.e., they become part of the *vehicle* for cognition, which is why this view is also sometimes called *vehicle cognition*). Note that our focus will be on factors external to the subject's brain and central nervous system that are specifically information-processing, such that the idea is that at least some of the information-processing that is part of the wider cognitive process is taking place outside the subject's brain and central nervous system.⁵

The standard criterion for extended cognition (the 'parity argument') is that an extended cognitive process is one that is functionally on a par with a comparable on-board cognitive process, in virtue of how it is seamlessly integrated into the subject's cognitive character, leading to rich feedback loops in its employment. Neuromedia clearly fits the bill on this score. Indeed, neuromedia seems to constitute a far more compelling example of extended cognition than the usual cases that are offered in the literature. The original example offered of an extended cognitive process is that of the dementia sufferer, 'Otto', who makes up for his failing memory by employing a notebook that he carries around with him.⁶ The thought is that if Otto always has the notebook with him, and regularly uses it to, say, navigate his environment, then this would constitute an extended cognitive process.

But this isn't all that plausible. For one thing, the technology does not seem especially seamless and integrated in its use, and the feedback loops are somewhat thin, given the limitations of the media in play. In particular, is Otto's use of the notebook really as seamless as his employment of his biological memory (when it is working anyway)? Relatedly, Otto's use of the notebook is arguably phenomenologically very different to his use of his biological memory.⁷ One experiences consulting one's notebook as using an external instrument; this is not the case when

we retrieve our biological memories. That is, Otto's relationship to the notebook is very much one of subject-and-instrument, unlike his use of his on-board cognitive resources.

Things look a bit better if we shift our attention from notebooks to wearable tech, like an Apple Watch. Since it's wearable, one can imagine that it does start to become seamlessly integrated into one's cognitive practices over time, perhaps to the point that one is no longer always conscious that one is employing it. The sophisticated nature of the technology also means that there is more scope for rich feedback loops on a number of fronts. Looking at the watch might provide information, which in turn stimulates biological memories, which in turn might influence what you do next, thereby generating new information, and so on. So we have one's on-board cognitive resources working in an integrated fashion with the technology.

I grant that a case like this could, over time, pass the test for being a case of extended cognition, in the sense of being functionally on a par with corresponding instances of on-board cognition. Neuromedia would be a much more convincing example of extended cognition, however. After all, even with wearable technology there is still an inevitable sense in which one's relationship to the technology is still, at least phenomenologically, one of subject-and-instrument, such that they are (quite often, anyway) not phenomenologically on a par. In contrast, it is built into the very idea of neuromedia that our interactions with the technology can be so seamless that we are not always even aware that we are employing technology, so it will be a much clearer-cut case of extended cognition than, say, an Apple Watch or a pair of Google Glasses.⁸

Neuromedia is also an interesting case to focus on because, unlike other potential forms of extended cognition, there is a sense in which it needn't be a technological adaption that is external to the skin and skull of the subject (it is thus in one sense an 'internal' form of extended cognition). In particular, if the neuromedia is developed along the cyborg route described above, then the technology may well be completely within the skin of the subject (although it will still be in a sense 'external' to the subject's brain and central nervous system, even if it is somehow embedded within it). That the technology is 'hidden' in this way is of course part of what helps to make one's use of it so seamless, as one is not physically interacting with the technology at all.

Note that even if one grants the possibility of extended cognition, it is a further question whether such cognition generates epistemic states like knowledge (i.e., whether there is *extended knowledge*). Perhaps only the subject's on-board cognitive processes are able to generate the epistemic pedigree required for knowledge. I've argued elsewhere that on the most plausible accounts of the nature of knowledge—roughly, *virtue epistemology*, broadly conceived—there is no inherent reason to be sceptical about extended knowledge.⁹

According to virtue epistemology, a necessary condition on knowledge is that one's cognitive success is significantly attributable to one's manifestation of cognitive agency.¹⁰ The crux

of the matter is that so long as the extended cognitive process is suitably integrated within the subject's cognitive character (which it needs to be, if it is to count as an extended cognitive process in the first place), then it will count as one of the subject's extended cognitive processes. Thus any target cognitive success that results from this process will be significantly attributable to the extended cognitive subject (i.e., her cognitive character as a whole, including the extended cognitive process). This would be extended knowledge.¹¹

One question we might ask about extended knowledge is whether it is in some sense second-grade knowledge when compared with unextended knowledge. One rationale for this is that we are epistemically dependent upon the technology for our knowledge, rather than being self-reliant. There is certainly something to this thought, in that our reliance on technology brings with it a kind of epistemic vulnerability, in that we run the risk of being cognitively impoverished were the technology to fail us (such is the premise of many a disaster movie). But the mere fact that we are dependent on technology for our knowledge is not itself an obvious worry. In fact, it is quite common for our knowledge to be dependent upon external factors that, alongside our (unextended) cognitive abilities, play a significant explanatory role in our cognitive success. This is what I have elsewhere called *epistemic dependence*.¹² Nonetheless, there is an important issue here, concerned with the distinctive value of certain epistemic traits, to which we will return.

4. THE EPISTEMIC AIMS OF EDUCATION

Education has many goals, some of them *social* (e.g., to help students to get along with each other), some of them *practical* (e.g., to enable students to have certain useful skills, like good handwriting), some of them *political* (e.g., creating good citizens), and so on. But one core goal of education is specifically *epistemic*. That is, we want to confer epistemic skills and states onto students, so that they are good inquirers, know useful information, can reason well, and so on.

I've argued elsewhere that the overarching epistemic goal of education should be to promote *intellectual character* rather than bodies of information/knowledge or the development of mere cognitive skills. In particular, intellectual character essentially involves the development of a specific kind of cognitive skill: *intellectual virtue*.¹³ (Related to this, I also claim that education should be geared towards promoting understanding rather than mere rote knowledge. We will come back to this point in due course).

Intellectual virtues—such as *conscientiousness* or *open-mindedness*—are distinct from mere cognitive skills along several axes.¹⁴ The list of differences is in fact very long, but let's focus on some of the key divergences. To begin with, intellectual virtues involve distinctive *motivational states*

that are constitutive of possessing the virtue. In general terms, the motivational state associated with an intellectual virtue is a love of the truth. But such motivational states are not a constitutive part of mere cognitive abilities. Indeed, one can manifest a cognitive ability such as one's rational faculties even while having no particular concern for the truth. Perhaps, for example, one is a lawyer who simply wants to develop a strong case for the prosecution, even though one's personal view is that the defendant is innocent, and one simply doesn't care what the truth of the matter is, but only what kind of compelling case can be presented the court. In this case one is manifesting a high level of cognitive ability, but one is not manifesting intellectual virtue, as the cognitive ability is reflecting one's instrumental goals rather than a love of the truth.¹⁵

Like other virtues (arguably, at any rate), intellectual virtues have the property of *lying between two vices*; one of excess and one of deficiency. One can be lacking in conscientiousness (deficiency), in which case one lacks this intellectual virtue. But one can also be overly conscientious (excess), perhaps by obsessively attending to every detail, no matter how trivial, in which case one will also lack this intellectual virtue. Mere cognitive abilities are not like this. One's perceptual faculties may be very reliable, but they can always be more reliable, and if they are then this is a good thing. For cognitive abilities the general rule is that we evaluate them in terms of their reliability, and the more reliable they are the better. If that's right, then there is no vice of excess when it comes to cognitive ability.

Intellectual virtues have a *distinctive kind of value* that contrasts them with mere cognitive abilities. They are to be prized regardless of their practical worth, for example, whereas mere cognitive abilities are usually only evaluated in terms of whether they serve our instrumental goals. Relatedly, intellectual virtues are also held to be constitutive parts of a life of flourishing, and thereby have a kind of non-instrumental, final, value (on account of the fact that the life of flourishing is meant to be valuable for its own sake). In contrast, mere cognitive abilities are axiologically evaluated in terms of how useful they are. For example, if a mere cognitive ability is no longer practically useful, then there would be nothing intellectually amiss in a subject choosing to let this ability wane. Crucially, however, this is not so of an intellectual virtue, in that the intellectually good person would recognise the inherent value of an intellectual virtue and strive to maintain it.

Following on from this point, intellectual virtues are distinct from mere cognitive abilities in terms of their *acquisition* and *cultivation*. Mere cognitive abilities can be innate, such as our cognitive faculties like memory, and they can also be acquired unreflectively, such as by one being continually exposed to the relevant stimuli. But intellectual virtues, like virtues more generally, are not like that. They need to be acquired in a *reflective* fashion, such as by emulation of someone who already has the intellectual virtues. Moreover, once acquired one needs to cultivate one's

intellectual virtues, since if they are not cultivated then they are lost (again, this is in contrast to mere cognitive abilities, which are often not lost once acquired, even if not cultivated).

The intellectual virtues promote *epistemic autonomy*, where this is the ability to develop and cultivate one's own viewpoints on matters of interest, and to determine what matters in the first place. One's mere cognitive abilities may enable such epistemic autonomy, but this is more like a side-effect rather than being central to the kind of abilities that they are. Relatedly, the intellectual virtues promote *active knowing* where this involves an inquiry/curiosity-driven approach to knowledge and understanding. This is contrast to *passive knowing*, where one's knowledge is simply passively received, and is not the product of thirst for it, of a seeking out of it.¹⁶

This feature of the intellectual virtues also explains why an account of the epistemic aims of education in terms of the development of intellectual character dovetails with a view according to which education should promote understanding rather than mere fact/skill retention, or rote knowledge. This is because understanding is by its nature an active epistemic state, unlike mere knowledge, which can be passively acquired. One can know something just by being told it, for example. But to understand something—a mathematical principle, say—it's not enough that one truly believes it via a good epistemic source (such as testimony from an expert). Rather, one needs to be grasp why this principle is true, how the different aspects of the principle relate to one another. Moreover, one needs to be able to employ this principle appropriately, where this means more than just asserting it when asked to do so. That is, one manifests one's understanding by being able to do things with what one's understands. All these features make understanding essentially an active epistemic standing, which is why developing intellectual character naturally leads to creating inquirers who seek to understand, and not merely know.¹⁷

Finally, the intellectual virtues perform a *managerial role* in one's cognitive architecture, in that they are employed to govern the use of cognitive skills rather than *vice versa*. One's intellectual virtues will determine what one cares about from an epistemic point of view, and thus determine the nature of the inquiries that one undertakes. In this way the intellectual virtues will marshal one's other epistemic resources, such as one's mere cognitive skills and one's body of knowledge, to promoting these ends. One consequence of this feature of intellectual virtues is that they have a more general focus than mere cognitive abilities. The latter are usually, if not always, an ability to do something quite specific, whereas intellectual virtues are general capacities that can be implemented in lots of distinct ways (think, for example, of the multiple ways in which being intellectually conscientious can manifest itself).

I think that once we understand how the intellectual virtues function in our cognitive lives, and their inherent value, then the claim that the epistemic goal of education ought to be the development of intellectual character ceases to be controversial. There may be facts and mere

cognitive skills that every person ought to have, in which case education should instil them. But this will be just a starting-point in the educational process. What we really want is to develop students' intellectual character so that they have a body of intellectual virtues that enables them to inquire well, and thereby to employ their mere cognitive skills and body of knowledge to intellectually valuable ends. There is so much more, from an epistemic point of view, to education than merely instilling skills and facts into the subject.¹⁸

5. NEUROMEDIA AND THE EPISTEMOLOGY OF EDUCATION

How will neuromedia affect our educational practices? One effect is that it will make a lot of our current educational practices redundant. Knowledge and skills that might hitherto have needed to be taught can be technologically engineered instead. Whereas students might have previously acquired knowledge using their own on-board cognitive resources, they will now be able to draw on their extended cognitive processes too. Why would we want students to remember large bodies of factual information using their biological memory, when they can draw on their extended memory and its vast resources? The same goes for lots of our basic cognitive skills. Why teach people to learn a foreign language if anyone can, via the technology, simply speak any language they want immediately?

It thus seems that there is a lot less for the educator to do in a world where neuromedia is common, and that's because many of our cognitive processes can be off-loaded onto the new technology. So one might well wonder what the epistemic point of education would be in such a scenario. In particular, is there anything left for the educator to do, or can the whole educational enterprise be off-loaded onto the technology? Crucially, however, while our cognitive abilities can be extended via neuromedia, the idea that our intellectual virtues can be extended in this fashion is somewhat implausible. If that's right, then while neuromedia will drastically reduce the need to educate people for basic knowledge and cognitive skills, it's primary epistemic function—that of developing intellectual virtue, and thus intellectual character—will remain (albeit in a slightly altered form, for reasons that I will explain in the next section).

In order to see why the intellectual virtues cannot be simply off-loaded onto technology like other kinds of cognitive ability, let's look at an intellectual virtue alongside a comparable (mere) cognitive ability. For example, let's consider the intellectual virtue of being observant with the mere cognitive ability involved in having good perceptual abilities. For instance, we could

contrast the exceptional observational skills manifested by Sherlock Holmes with the merely good perceptual abilities of his sidekick Watson.

Being observant involves so much more than merely having good perceptual faculties, as the contrast between Holmes and Watson makes clear, as there is such an epistemic gulf between them on this score. Faced with the very same visual scene, for example, Holmes is able to extract vast quantities of useful information, unlike Watson. This reflects the fact that perception is just a cognitive ability, but that being observant is an intellectual virtue. As such, the two traits differ along the axes noted above.

Perception is often passive, for example, but being observant is essentially active, in that one is actively inspecting the scene before one for information. As such being observant reflects one's epistemic autonomy. Relatedly, being observant involves a motivation to find the truth, something that could be completely lacking in one who merely has good perceptual faculties. One is not born an observant person, but one needs to acquire and cultivate this trait, and that will require one to reflect upon one's exercise of it. In contrast, one can simply have good perceptual faculties, and may not need to do anything to ensure that they are retained. They can also be reliably exercised in a completely unreflective manner.

The intellectual virtue of being observant also lies between two vices, though this might not be immediately obvious. Doesn't Holmes have this trait in excess, and if so, doesn't that mean that he lacks this virtue? I don't think that's right, and this becomes clear once we reflect on what it would mean to be excessively observant. This is not to have acute observational skills like Holmes, but rather to obsess about irrelevant details. This is why Holmes, even while being exceptional in this regard, nonetheless still retains the virtue, as he is attending only to the details that matter, and not merely every possible detail, regardless of its import. In any case, we don't evaluate perceptual abilities in this way. If one's vision, say, is drastically improved, then that's always a good thing from an epistemic point of view. There's no such thing as having vision that's 'too good' to count as a genuine cognitive ability.

Notice too how while perception can very easily, and passively, lead to knowledge of one's environment, being observant enables one to actively understand things as a result of one's perception. Watson surveys the crime scene in front of him and immediately comes to know facts of various kinds, such as that the victim is such-and-such, that the window is open, that there is a revolver on the mantelpiece, and so on. But Holmes does not merely come to know all these facts, but also appreciates their significance and thereby comes to understand something—e.g., that this murder scene has been faked, that such-and-such cannot be the murderer, that such-and-such must be involved in this crime, and so on.

Finally, recall our point about how the intellectual virtues play a managerial role in our cognitive economy. Sherlock is putting his knowledge and other cognitive skills, such as his perceptual skills, at the service of his intellectual virtue. It is his intellectual character, composed of his intellectual virtues, that is determining the lines of inquiry that he takes, and which employs his vast body of knowledge and his exceptional cognitive skills. Notice that this is an essentially reflective process, just as the cultivation of one's intellectual virtues is, in that it involves deliberation and judgment. It is not the kind of thing that can 'just happen' (which is precisely how perception often occurs).

I don't think that it is controversial that many of Watson's perceptual abilities could become cognitively extended via neuromedia. Take his eyesight, for example. One could certainly imagine cognitive augmentations of this faculty that enabled him to do incredible things with his vision, to be able to see far into the distance, or focus in on a particular detail before him at high magnitude. Moreover, such an augmentation could, over time, be so integrated within his other cognitive abilities that his employment of it is completely seamless, such that it doesn't even feel like he is relying on technology at all, but rather just using 'his' cognitive abilities.

Could we do the same with Holmes' intellectual virtue of being observant? I don't see how. Holmes may well employ technology in lots of ways to assist his observational powers, as when he uses his famed magnifying glass to inspect the crime scene before him. But notice that this is technology that is being reflectively brought into service to serve his intellectual virtues. Similarly, even if Holmes is fitted with the same neuromedia to enhance his vision that we just hypothetically attributed to Watson, it is still only employed under the guidance of his intellectual virtue.

Indeed, what would it even *mean* for an intellectual virtue to be cognitively extended? One's intellectual virtues, and one's intellectual character more generally, are constituted by being reflective, managerial traits that guide one's employment of one's cognitive abilities (extended or otherwise) and one's knowledge (again, extended or otherwise). This means that it is built into these traits that they are manifestations of one's unextended cognitive character. Technology, even in the form of neuromedia, is only ever a tool that one can virtuously employ in reaching one's intellectual goals; it is never a substitute for the intellectual virtue itself.

If that's right, then while a great deal of 'lesser' educational tasks will not be necessary in a world of neuromedia, the core epistemic goal of education—that of developing good intellectual character, and thus promoting the intellectual virtues—will remain intact. Moreover, one could argue that in an age of neuromedia we need intellectual character more than ever. Consider some of the challenges that neuromedia poses that we looked at above, such as how it can be used just

as much to misinform as much as inform. The intellectually virtuous will be far better placed to handle such difficulties.

I think this point is also relevant when it comes to our increasing reliance on technology, something that we noted above seems problematic. The intellectually virtuous person will be reliant on technology, including neuromedia, in a very different way to someone who is lacking those virtues. For one thing, having intellectual virtues means having very general cognitive skills, as opposed to the way in which mere cognitive abilities tend, as we noted above, to be devoted to specific cognitive tasks. This means that they are very practically useful in terms of helping one to develop further mere cognitive skills, and in obtaining new knowledge. As a consequence, in a situation where the technology is suddenly no longer available, then the intellectually virtuous will be in a good position to know what to do, and in particular to adapt to the new epistemic environment.

In addition, the intellectually virtuous person is also someone who will be careful about their reliance on technology in the first place. After all, these extended cognitive processes are, for them, a mere epistemic resource, and it is part of what it is to be intellectually virtuous that one employs one's epistemic resources wisely. Accordingly, the intellectually virtuous person will make sure that they are not overly reliant on technology, and in particular that they have the means available to them to function adequately were the technology to no longer be available.

6. INTELLECTUAL CHARACTER AND TECHNOLOGY

It is important to note that denying that the intellectual virtues cannot be cognitively off-loaded onto technology does not entail that technology can't be used to enhance one's employment of the intellectual virtues, or even that it can't play a role in the acquisition and cultivation of the intellectual virtues. The first point is obvious. Sherlock Holmes, our exemplar of intellectual virtue (albeit an exemplar who is often somewhat lacking in some of the non-intellectual virtues), frequently uses technology to aide his employment of his intellectual virtue (whether it is the magnifying glass from the original novels, or the advanced tech found in contemporary presentations of the detective).

The second point—that technology can play a role in developing intellectual virtue, and thus intellectual character—is perhaps not so obvious. I want to spell-out what this might involve by describing: (i) an actual project that we undertook, using technology to develop intellectual character; and (ii) a possible app that would help students enhance their intellectual character. As

we will see, in both cases we have technology at the service of the goal of enhancing intellectual character, but in neither case is it leading to instances of extended cognition. Moreover, turning the technology in question into neuromedia won't make a difference.

The first example I want to focus upon is a project that I was involved with, and am still involved with, that brought philosophy—more specifically, critical thinking—into prison education. The goal of the project was to enhance the prisoners' intellectual character.¹⁹ To that end, they were aided by two factors. The first was an (off-line) version of the MOOC (= Massive Open Online Course) on 'Introduction to Philosophy' that we had created, along with supporting educational materials (e.g., handouts to go with each topic, a set of critical thinking problems, discussion topics, and so on).²⁰ The MOOC was designed to introduce people to philosophical topics without presupposing any philosophical background, and so was ideal for this educational setting. The second was a series of seminars using a particular way of teaching philosophy/critical thinking, known as *Community of Philosophical Inquiry*, or CoPI for short.²¹ This approach has been widely and effectively used in educational contexts, and it also has the advantage that, like the MOOC, it doesn't presuppose any previous knowledge of philosophy (which is why this technique is often used in philosophy in schools programmes that target younger children).

The project had a demonstrable effect on the prisoners' intellectual character. At the start of the project, the prisoners struggled to articulate their reasoning, struggled to understand other people's reasoning (or even grasp their different points of view), gave up on problems they found difficult very easily, were unwilling to collaborate with others in solving problems, showed very little creativity in problem-solving, and so on. These traits are all indications of a lack of intellectual character, and in particular showed that the prisoners—like many of us—have many intellectual vices. But on all these fronts the prisoners showed marked improvement as the project went on. For example, they became much better at articulating their reasoning and grasping the reasoning of others. This meant that they could now engage in a genuine reasons-based debate with each other, rather than simply dismissing each other's opinions from the off. Their intellectual tenacity improved, in that they stuck with difficult problems for longer, often engaging collaboratively with others in trying to solve the problem. Relatedly, they were more intellectually creative in their attempts to answer problems, employing novel approaches rather than sticking only to the same strategies. These are all indications of the development of intellectual virtue and thus of the development of intellectual character.²²

Note that I say that the project *developed* the prisoners' intellectual character and thus their intellectual virtues, which is not to say that they were suddenly intellectually virtuous after the project—that would have been an incredible educational feat, given that the project only lasted a

couple of months. The prisoners were instead just a lot closer to being intellectually virtuous as a result of the project, in that they had mastered some important intellectual skills, such as being able to articulate their reasoning, which are necessary to acquire intellectual virtue. More importantly, now that they had these skills they were in a position to further develop their intellectual character themselves.

The technology employed in this project—essentially, the MOOC—certainly enhanced the effectiveness of the educational techniques employed. Moreover, one could imagine running versions of this project which employed even more technology, including neuromedia. Perhaps neuromedia is used as a replacement for watching the MOOC, and instead of educators in the room running the CoPI sessions, one could imagine a kind of technologically ‘scaffolded’ environment that delivered prompts to the prisoners to help them master these techniques. But no matter how much of a cognitive role in this process is played by the technology, there is no way of simply technologically engineering the prisoners’ intellectual virtue. Even in the most technologically-enhanced version of this project one could imagine, one is at most using the technology to enable the development of certain basic intellectual skills that are crucial to a good intellectual character, such as learning more creative ways of problem-solving. Such skills are vital to the development of intellectual virtue, but they are not yet intellectual virtues, in that they are lacking many of the features of the intellectual virtues noted above. But as we saw above, turning them into intellectual virtues will require the subject manifesting all kinds of dispositions—e.g., distinctive motivational states, reflecting on one’s intellectual performance, actively rather than passively knowing, etc.,—and these dispositions are not of a kind that can simply be cognitively off-loaded onto technology. Instead, the manifestation of genuine intellectual virtue will always involve the technology becoming a mere instrument for the intellectually virtuous person to employ.

Now consider an idea for a possible educational app. Think of how one these days finds the answer to a question one has: we simply ‘Google it’. This puts a wealth of information right at our fingertips, and if neuromedia becomes a reality then it will be closer still (i.e., phenomenologically akin to our biological memory). But the intellectually virtuous person is not satisfied with merely getting an answer to a question. Rather, she will engage in intellectually virtuous *inquiries*, where this will involve an interesting series of questions being answered. Imagine an app that tries to cultivate this trait within a student, such that the student doesn’t merely ask a question of Google and get an answer, but thereafter follows-through with a train of inquiry. How might it be constructed?

The problem, of course, is how to have a way of determining a good inquiry in advance of having a fixed account of what constitutes a good inquiry in this context. But there are ways of resolving this difficulty. Suppose one starts with an educational pilot, whereby pupils are asked to develop a chain of inquiry based on an initial question, and are graded on how intellectually stimulating this inquiry is. This is obviously a thin data set, but nonetheless one could use this data to start ‘weighting’ certain inquiries over others, in terms of whether the person doing the inquiry has been judged to have done good inquiries in the past. If one could get the relevant algorithm right, then over time by repeating this process with more and more pupils and educators, one could get a way of rating inquiries. Pupils could be given a score as an inquirer, depending on what kind of inquiry route that they took, so that it is not just inquiry paths that are rated but also inquirers. In the beginning, this would involve lots of evaluation from the educators, but over time this would become a self-regulating process (at least if the algorithms are done well), in that one could have ‘good’ inquirers rate the inquiries of others (perhaps anonymized to ensure that there is no bias entering the system). Over time, one could develop a data set for particular questions that picked out lines of inquiry that were distinctive of good inquirers, and which also picked out good inquirers too.

With this data set in play, such an app could go beyond the developmental stage and actually play a role in the cognitive development of pupils, by offering a way of helping them to enhance their core critical skills involved with critical inquiry, skills that are essential to the development of intellectual virtue. Note that expanding the initial run to a broader set of pupils also substantially expands the data set, and thereby helps to refine the algorithm in place to determine good inquiries, and thus good inquirers. We now have a bigger set of pupils rated by their educators as being good inquirers (because their inquiries have been rated as good), and we now have data about how inquirers are rating each other (where this is weighted in terms of how good an inquirer one is).

Suppose that one now takes this basic model, gleaned from the specifically educational context, and makes it available wholesale. One could imagine, for example, that instead of doing a simple Google search, one has the option of doing instead a ‘Schmoogle search’ (or whatever it might be called). This would involve getting to see where ‘good’ inquirers went next in terms of their inquiries originating with this search. This way one would get to see what good inquirers did next, as opposed to simply getting a closed answer to a closed question.

More interestingly, however, moving the model to a wholesale scale affords a number of additional advantages. Before we had the educators as expert, with a weighed rating as a result. Imagine now that inquirers are rating each other *en masse*. An algorithm needs to be developed to ensure that raters are weighted accordingly (e.g., in terms of their past inquiries and how they fit

with previously judged to be ‘good’ inquiries, and in terms of whether their ratings of inquiries fits with the ratings from ‘good’ evaluators). With the right algorithms in place, there will be a way of inquirers ranking others and being ranked themselves that promotes good inquiry. One can even imagine that, over time, inquirers will not merely Google search, but will actively ‘Schmoogle search’ instead, and even that inquirers who are rated as ‘good’ via this process are lauded, with their particular searches followed by others.²³

This is all highly hypothetical, of course, but the point of the exercise is to demonstrate how technology could be employed to help people enhance their intellectual character. As with the prisoners, however, all that is on offer here is an *enabler* of the development of intellectual character—the technology cannot all by itself ensure that pupils develop the suite of intellectual virtues distinctive of a good intellectual character. What is being developed are some basic cognitive skills that are necessary for the development of cognitive character. But for these cognitive skills to develop into intellectual virtues the subject is going to have to manifest the dispositions listed above—e.g., being suitably reflective, being guided by the right kind of motivational states, being an active rather than passive knower, and so on—that can’t be off-loaded onto the technology. Again, we see that there is a place for technology to aid the development of intellectual character, but it cannot be a substitute for it.

7. CONCLUDING REMARKS

Were neuromedia to become a reality, much of our knowledge would likely be replaced by extended knowing, and many of our cognitive skills would likely end up being extended cognitive skills. But this doesn’t pose any challenge to the epistemic goal of education, since this was never about the development of (mere) cognitive skills or the instillation of (mere) knowledge. Instead, its epistemic goal is that of developing intellectual character, and thus the intellectual virtues that comprise a good intellectual character. This goal would be no less important in an age of neuromedia, because, as we have seen, intellectual virtues are by their nature non-extended cognitive traits. If anything, it is *more* important to inculcate intellectual character in an age of neuromedia. In any case, while technology might have a role to play in aiding the development of intellectual character, it cannot play any more of a role than that.²⁴

REFERENCES

- Adams, F., & Aizawa, K. (2008). *The Bounds of Cognition*, Blackwell, Oxford.
- Ahlstrom, A. (2013). *Epistemic Paternalism: A Defence*, London: Palgrave Macmillan.
- Axtell, G. (1997). 'Recent Work in Virtue Epistemology', *American Philosophical Quarterly* 34, 410-30.
- Bachr, J. (Ed.) (2015). *Intellectual Virtues and Education: Essays in Applied Virtue Epistemology*, London: Routledge.
- (2011). *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*, Oxford: Oxford University Press.
- Batttaly, H. (2014). 'Intellectual Virtues', *Handbook of Virtue Ethics*, (ed.) S. van Hooft, 177-87, London: Acumen.
- Bovill, M., & Pritchard, D. H. (2015). 'Philosophy in Prisons: Critical Thinking and Community of Philosophical Inquiry', *unpublished manuscript*.
- Chemero, T. (2009). *Radical Embodied Cognitive Science*, Cambridge, MA: MIT Press.
- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*, Oxford: Oxford University Press.
- Clark, A., & Chalmers, D. (1998). 'The Extended Mind', *Analysis* 58, 10-23.
- Goldman, A. (1991). 'Epistemic Paternalism: Communication Control in Law and Society', *Journal of Philosophy* 88, 113-31.
- Greco, J. (2003). 'Knowledge as Credit for True Belief', *Intellectual Virtue: Perspectives from Ethics and Epistemology*, (eds.) M. DePaul & L. Zagzebski, 111-34, Oxford: Oxford University Press.
- (2007). 'The Nature of Ability and the Purpose of Knowledge', *Philosophical Issues* 17, 57-69.
- (2008). 'What's Wrong With Contextualism?', *Philosophical Quarterly* 58, 416-36.
- (2009). *Achieving Knowledge*, Cambridge: Cambridge University Press.
- (2012). 'A (Different) Virtue Epistemology', *Philosophy and Phenomenological Research* 85, 1-26.
- Hardwig, J. (1985). 'Epistemic Dependence', *Journal of Philosophy* 82, 335-49.
- Hutchins, E. (1995). *Cognition in the Wild*, Cambridge, MA: MIT Press.
- Lipman, M. (1991). *Thinking in Education*, New York: Cambridge University Press.
- Lynch, M. P. (2014). 'Neuromedia, Extended Knowledge, and Understanding', *Philosophical Issues* 24, 299-313.
- (2016). *The Internet of Us: Knowing More and Understanding Less in the Age of Big Data*, New York: W. W. Norton & Co.
- McCall, C. (2009). *Transforming Thinking: Philosophical Inquiry in the Primary and Secondary Classroom*, London: Routledge.
- Noë, A. (2004). *Action in Perception*, Cambridge, MA: MIT Press.
- Palermos, S. O. (2014a). 'Knowledge and Cognitive Integration', *Synthese* 191, 1931-51.
- (2014b). 'Loops, Constitution, and Cognitive Extension', *Cognitive Systems Research* 27, 25-41.
- Palermos, S. O., & Pritchard, D. H. (2016). 'The Distribution of Epistemic Agency', *Social Epistemology and Epistemic Agency: De-Centralizing Epistemic Agency*, (ed.) P. Reider, 109-26, Lanham, MD: Rowman & Littlefield.
- Pritchard, D. H. (2018). 'Extended Knowledge', *Extended Epistemology*, (eds.) J. A. Carter, A. Clark, J. Kallestrup, S. O. Palermos & D. H. Pritchard, 90-104, Oxford: Oxford University Press.
- (2017). 'Extended Virtue Epistemology', *Inquiry* [Online First, DOI: 10.1080/0020174X.2017.1355842].
- (2016a). 'Epistemic Dependence', *Philosophical Issues*, 30, 1-20.
- (2016b). 'Epistemic Risk', *Journal of Philosophy* 113, 550-71.
- (2016c). 'Intellectual Virtue, Extended Cognition, and the Epistemology of Education', *Intellectual Virtues and Education: Essays In Applied Virtue Epistemology*, (ed.) J. Bachr, London: Routledge.
- (2016d). 'Seeing It For Oneself: Perceptual Knowledge, Understanding, and Intellectual Autonomy', *Episteme* 13, 29-42.

- (2014a). 'Knowledge and Understanding', *Virtue Scientia: Bridges Between Virtue Epistemology and Philosophy of Science*, (ed.) A. Fairweather, 315-28, Dordrecht, Holland: Springer.
- (2014b). 'Virtue Epistemology, Extended Cognition, and the Epistemology of Education', *Universitas: Monthly Review of Philosophy and Culture* 478, 47-66.
- (2013). 'Epistemic Paternalism and Epistemic Value', *Philosophical Inquiries* 1, 1-37.
- (2013). 'Epistemic Virtue and the Epistemology of Education', *Journal of Philosophy of Education* 47, 236-47.
- (2012). 'Anti-Luck Virtue Epistemology', *Journal of Philosophy* 109, 247-79.
- (2010). 'Cognitive Ability and the Extended Cognition Thesis', *Synthese* 175, 133-51.
- (2009). 'Knowledge, Understanding and Epistemic Value', *Epistemology (Royal Institute of Philosophy Lectures)*, (ed.) A. O'Hear, 19-43, Cambridge: Cambridge University Press.
- Pritchard, D. H., Millar, A., & Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations*, Oxford: Oxford University Press.
- Rowlands, M. (2009). *The Body in Mind: Understanding Cognitive Processes*, New York: Cambridge University Press.
- Rupert, R. D. (2009). *Cognitive Systems and the Extended Mind*, Oxford: Oxford University Press.
- Shapiro, L. (2011). *Embodied Cognition*, New York: Routledge.
- Sosa, E. (1991). *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge: Cambridge University Press.
- (1999). 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13, 141-54.
- (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford: Oxford University Press.
- (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge*, Oxford: Oxford University Press.
- (2015). *Judgment and Agency*, 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.
- (1999). 'What is Knowledge?', *Blackwell Guide to Epistemology*, (eds.) J. Greco & E. Sosa, 92-116, Oxford: Blackwell.

NOTES

¹ For more on neuromedia in a philosophical context, see Lynch (2014; 2016), though since it isn't clear to me that we are using this terminology in quite the same way, the reader should focus on the account of this notion offered here. I also encountered this terminology at a fascinating 'impact' workshop that we held at the University of Edinburgh's *Eidyn* research centre in 2015, as part of the AHRC-funded 'Extended Knowledge' project that I was running at the time. This brought academics from several disciplines together with those involved in the development of new technologies at companies such as Microsoft, IBM, and Google. Intriguingly, the consensus among the tech-gurus was that neuromedia was imminent (some even thought that it already existed, at least in prototype form), though there was disagreement about what exact form it would take. On this latter point, see also endnote 2.

² I was made aware of this distinction regarding two competing technological models for neuromedia, and the associated terminology to describe it, by speaking to technology industry delegates at the event described in endnote 1.

³ Of course, the rich already have cognitive advantages over the poor in terms of such things as their access to good education and cognitively useful technology, so this would be an *additional* layer of cognitive privilege.

⁴ The *locus classicus* for discussions of extended cognition is Clark & Chalmers (1998), but see also Clark (2008).

⁵ For two prominent critiques of extended cognition, see Adams & Aizawa (2008) and Rupert (2009). Note that extended cognition can also have a social aspect, as when the extended cognitive process is socially distributed. I will set this aside in what follows here, but have explored the particular features of this kind of extended cognition in Palermos & Pritchard (2016). See also Hutchins (1995), which is a seminal work on socially distributed cognition. Note too that the way I have described extended cognition treats embodied cognition as a form of extended cognition (sometimes the former is defined in such a way as to distinguish it from the latter). This is when features connected to one's embodiment play a constitutive role in one's cognitive processes. Again, I will be setting this type of extended cognition to one side for our current purposes, but for further discussion of embodied cognition, see Noë (2004), Chemero (2009), Rowlands (2009), and Shapiro (2011).

⁶ This case appears in Clark & Chalmers (1998).

⁷ Proponents of extended cognition are usually happy to grant that there are phenomenological differences of this kind, though they might also insist that they are overstated. For them the focus is a functional equivalence between the extended cognitive process and a corresponding non-extended cognitive process, and that's held to be compatible with there being some differences at the level of phenomenology. In any case, what everyone will surely agree with is that cases of extended cognition which involve both functional equivalence and a similar kind of phenomenology would be more compelling cases of extended cognition.

⁸ Note that I am not saying that there is any difference in kind here between neuromedia and other putative cases of extended cognition, but only that there is a difference in degree. For example, there will likely be cases where one is aware that one is employing technology even in the neuromedia case, as when one finds oneself 'remembering' something that one is aware one can't possibly have known in this way (e.g., an extremely complex mathematical theorem). Of course, neuromedia might develop in such a way that subjects do not become aware that they are using technology in such cases (e.g., in terms of the case just considered, perhaps they are also 'fed' information to suggest that they are mathematically adept), but remember that our interest is in neuromedia as a *bona fide* cognitive augmentation rather than as something epistemically malign. There are interesting issues here about whether certain kinds of epistemic paternalism, of a kind that would mirror the scenario just described whereby subjects are 'fed' faulty information for supposedly their own epistemic good, are permissible. For more on the nature and justification for epistemic paternalism, see Goldman (1991), Ahlstrom (2013), and Pritchard (2013).

⁹ See Pritchard (2010; 2017; 2018).

¹⁰ Indeed, some virtue epistemologists—e.g., Sosa (1991; 2007; 2009; 2015), Zagzebski (1996; 1999), and Greco (2003; 2007; 2008; 2009; 2012)—think that a condition of this kind is in fact sufficient for knowledge, a position I have elsewhere christened as 'strong', or 'robust' virtue epistemology (e.g., Pritchard, Millar & Haddock (2010, chs. 1–4)). My own view, however, is that it is not sufficient, in that we also need a distinct way of eliminating certain kinds of epistemic luck/risk (e.g., Pritchard 2012; 2016a). In any case, since the necessity claim is the more liberal of the two, in that it captures a broader range of virtue-theoretic proposals about knowledge, it will function very well for our purposes.

¹¹ For more on this point, see Pritchard (2010; 2017; 2018).

¹² See Pritchard (2016a). Note that this is very much a term of art, and so what I mean by it is not what, for example, Hardwig (1985) means by it in his influential article on epistemic dependence (though inevitably there are some overlaps).

¹³ See Pritchard (2013; 2014b; 2016b),

¹⁴ For more on the intellectual virtues, see Zagzebski (1996), Baehr (2011), and Battaly (2014).

¹⁵ Most real-life cases of legal advocacy are likely to be more complex on this front, in that even while the lawyer might be advocating for a viewpoint that they don't personally endorse, nonetheless they are manifesting intellectual virtue. That is, one can care about the truth, but also be willing to advocate for something other than what one knows is the truth for legitimate non-epistemic reasons, as might be the case in a legal context where one's ultimate concern is a free trial. The example just given is not of this kind, however, in that the lawyer is by stipulation only ultimately motivated by purely strategic, rather than epistemic, concerns.

¹⁶ For more on the importance of active knowing in the context of epistemic autonomy, see Pritchard (2016*d*).

¹⁷ For more discussion of this point about the difference between knowledge and understanding, see Pritchard (2009; 2014*a*) and Pritchard, Millar & Haddock (2010, ch. 4).

¹⁸ A related issue in this regard is that on most conceptions of the virtues, one cannot develop a subject's intellectual virtues without thereby developing their moral and practical virtues too, given how inter-related they all are (on this score, see also our prisons education project, discussed below, and also in endnote 22, which while focused on the development of prisoners' intellectual character, also developed their character more generally). This means that the epistemic goal of education in developing intellectual character would go hand-in-hand with a conception of the broader goal of education to develop the subject's virtuous character in general.

¹⁹ This project was hosted by the University of Edinburgh's *Eidyn* research centre, but was a collaboration between *Eidyn*, the Moray House School of Education (also at the University of Edinburgh), colleagues at New College Lanarkshire involved in prison education, and the Scottish Prison Service. Two prisons were targeted, a male prison (Low Moss) and a female prison (Cornton Vale). A second run of this project has just been completed, with more runs planned for the future. For more details about the project, see <http://eidyn.ppls.ed.ac.uk/project/philosophy-prisons>.

²⁰ Note, by the way, that the reason why the MOOC they used was off-line is that in the Scottish system prisoners are not allowed access to the internet. The version they used was nonetheless functionally equivalent to the on-line version, albeit without the online discussion forums. Incidentally, the MOOC in question has been enormously popular, with well over 2M enrolments worldwide. For more details about this course, go to: <https://www.coursera.org/learn/philosophy>.

²¹ The particular version of CoPI that we employed was pioneered by Catherine McCall, who also advised us on the project. See McCall (2009) for details of this approach. See also earlier work by Lipman (1991), which is a precursor to this model.

²² The results from the project, and their evidential basis, are detailed in Bovill & Pritchard (2015). A summary of the project and its results is available here: <https://www.ed.ac.uk/education/rke/making-a-difference/philosophy-in-prisons>. Note that one of the results of the project, which on reflection is unsurprising given standard conceptions of the virtues, is that in developing the prisoners' intellectual character we also seemed to develop their character more generally (e.g., their self-esteem, their respect and concern for others, and so on).

²³ Indeed, one could even envisage that, over time, the app becomes more nuanced in its evaluations, such that inquiries are evaluated along several axes (e.g., originality, novelty, etc.).

²⁴ An earlier version of this paper was presented at a conference on the Philosophy of Education at the University of London in May 2017, and I am grateful to the audience for their feedback on the talk. Thanks also to Mary Bovill, J. Adam Carter, Michel Croce, Andy Clark, Chris Kelp, Andrea English, Aaron James, Catherine McCall, Orestis Palermos, John Ravenscroft, and especially Michael Lynch. Thanks too to an anonymous referee from *Metaphilosophy* and David Mott from IBM's Emerging Technology team. This paper has benefitted from four grants, all of them for projects hosted at the University of Edinburgh's *Eidyn* research centre. The first is the AHRC-funded 'Extended Knowledge' project (AH/J011908/1). The other three are all funded by the John Templeton foundation. These are: (i) the 'Virtue Epistemology, Epistemic Dependence and Intellectual Humility' project, which was itself part of the wider 'Philosophy and Theology of Intellectual Humility Project' hosted by Saint Louis University; (ii) the 'Intellectual Humility MOOC' project; and (iii) the 'Philosophy, Science and Religion Online' project.