THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Extended cognition and humility

Citation for published version:

Pritchard, D 2020, Extended cognition and humility. in M Alfano, MP Lynch & A Tanesini (eds), *The Routledge Handbook of Philosophy of Humility*. Routledge Handbooks in Philosophy, Routledge.
<https://doi.org/10.4324/9781351107532>

Digital Object Identifier (DOI):

[10.4324/9781351107532](https://doi.org/10.4324/9781351107532)

Link:

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

Document Version:

Peer reviewed version

Published In:

The Routledge Handbook of Philosophy of Humility

Publisher Rights Statement:

This is an Accepted Manuscript of a book chapter published by Routledge in The Routledge Handbook of Philosophy of Humility on 7 July 2020, available online: <https://www.routledge.com/The-Routledge-Handbook-of-Philosophy-of-Humility/Alfano-Lynch-Tanesini/p/book/9780815364115>

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.

EXTENDED COGNITION AND HUMILITY

Duncan Pritchard

University of California, Irvine & University of Edinburgh

ABSTRACT. Extended cognition is when cognitive processes extend beyond the brain and nervous system of the subject, and in the process properly include such ‘external’ devices as technology. This paper explores what relevance extended cognitive processes might have for humility, and especially for the specifically cognitive aspect of humility—*viz.*, intellectual humility. As regards humility in general, it is argued that there are no in principle barriers to extended cognitive processes helping to enable the development and manifestation of this character trait, but that there may be limitations to the extent to which one’s manifestation of humility can be dependent upon these processes, at least insofar as we follow orthodoxy and treat humility as a virtue. As regards the cognitive trait of intellectual humility in particular, the question becomes whether this can itself be an extended cognitive process. It is argued that this wouldn’t be a plausible conception of intellectual humility, at least insofar as we treat intellectual humility (like humility in general) as a virtue.

1. EXTENDED COGNITION

Extended cognition is the thesis that a subject’s cognitive processes can extend beyond the brain and central nervous system of the subject; indeed, can extend beyond her skin and skull. In particular, it is the thesis that features of the subject’s cognitive environment, such as technology, can in the right conditions become genuine proper parts of the subject’s cognitive processes.¹ Extended cognition has in recent years become a very influential research programme in the cognitive sciences, with its philosophical implications developed and explored by philosophers of mind and cognitive science.

Putative examples of extended cognition are legion, and uniformly controversial. The canonical example, due to Andy Clark and David Chalmers (1998), concerns a subject Otto who is

losing his memory and so begins to use a notebook to compensate for his memory loss. The question is whether Otto's use of the notebook can eventually count as an extended cognitive process—i.e., as a kind of extended memory, on a par with his biological memory. Clark and Chalmers claim that it can. In particular, they appeal to the following parity principle:

“If, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is [...] part of the cognitive process.” (Clark & Chalmers 1998, 29)

Applied to Otto and his notebook, the rationale in play is that so long as the notebook functions just like his ‘internal’ onboard memory, then we should treat it as a genuine part of an extended cognitive process. That is, Otto's use of the notebook is not to be understood as simply an agent employing an instrument, but rather as an integrated part of his cognitive processes, and thus cognitively equivalent to his (biological) memory.²

The crucial element in this thought experiment is, of course, whether Otto's use of the notebook could ever be on a functional par with his use of his biological memory. In particular, the latter tends to have an immediacy and a distinctive kind of phenomenology that goes along with its usage that seems very different to how one might employ a notebook. For example, one can employ one's memory in one's reasoning in such a seamless fashion that one might not even be consciously aware that this is what one is doing. Could one's use of a notebook ever be seamless in this fashion? In particular, won't the use of the notebook always involve a kind of intellectual distance of just the kind present when one uses an instrument?

Still, even if one is not convinced that the parity principle is satisfied in this case, one might think that there are cases that are more plausible in this regard. Indeed, for the purposes of this article we can just stipulate a kind of technology that would fit the bill, and leave it an open question whether such technology yet exists. Elsewhere I have referred to such technology as *neuromedia*, which I described as follows:

“[...] information processing technology that is so seamlessly integrated with one's 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.” (Pritchard 2018d, 328)³

So understood, neuromedia would clearly satisfy the parity principle. Moreover, even if one is not convinced that there is any technology available just now that would count as neuromedia in this specific sense, I think it is undeniable that those who are developing new technologies of the

relevant kind are aiming to create something along just these lines. For example, couldn't we easily imagine a time when consulting the internet for information is so automatic and phenomenologically immediate that it is no different from consulting one's memory (to the extent that one might not ordinarily be able to tell the difference)?⁴ If so, then this would be neuromedia, and hence a plausible case of extended cognition, in virtue of its satisfaction of the parity principle.⁵

Notice too that neuromedia would also clearly manifest a further feature that is often thought to be entailed by the parity principle, which is that cognitively extended processes need to be sufficiently integrated into one's wider cognitive system that their employment leads to rich feedback loops. The information one gains from one's neuromedia could mesh in substantive ways with information from other sources (memorial, perceptual, and so on), such that they collectively guide action, leading in turn to new information-processing that employs both extended and onboard cognitive resources. This kind of cognitive integration is important to any claim about functional equivalence between the extended and the corresponding onboard cognitive processes, given that one's onboard cognitive processes are clearly cognitively integrated. Think, for example, about how one's feeling of coldness meshes with one's memorial and perceptual knowledge of one's environment (seeing that it is night time, remembering how cold it gets at night in these parts, and so on), with each reinforcing the other and guiding action accordingly (seeking out one's jumper, for example).⁶

Of course, one might dispute the parity principle, not just in terms of the detail (e.g., whether it entails cognitive integration, and to what extent) but more broadly as a test for extended cognition. Indeed, one might be suspicious of the very idea that there can be extended cognition, and hence the possibility of neuromedia would be neither here nor there.⁷ For our purposes, however, we will take it as given that there could be such a thing as extended cognition, and that were it to exist then neuromedia would be an instantiation of it. The question we will be engaging with is what implications the extended cognition thesis might have for humility.

2. COGNITIVELY EXTENDED HUMILITY

One challenge facing any discussion of the relationship between extended cognition and humility is that the former, but not the latter, specifically concerns cognitive processes. There are many conceptions of humility in the literature, but while they all involve a cognitive element (which we

will come to in a moment), none of them consider humility to be exclusively a cognitive process. Indeed, it is standard in this regard to distinguish between humility in general and a particular type of humility that *is* geared towards cognitive ends—*viz.*, *intellectual humility*. If we think of humility, as is common, as a virtue, and thus treat intellectual humility as a specifically intellectual virtue, then we capture this point by noting that while the general manifestation of humility involves a motivational state directed towards the good, the manifestation of intellectual humility involves the specific motivation towards the intellectual good—i.e., truth.⁸

Nonetheless, we will begin by considering what bearing extended cognition might have for humility in general (we will be considering intellectual humility in its own right in the next section). Even though virtues in general, as opposed to intellectual virtues, are not specifically cognitive traits, they do employ cognitive traits. Being virtuous is not simply a matter of wanting to be a certain way (having good motives, and so on), but rather also involves being reliable at attaining certain virtuous outcomes, at least in the right conditions. Accordingly, there are cognitive skills involved in manifesting virtues even if the virtues themselves are not geared towards specifically cognitive goals. With this in mind, we can ask whether the manifestation of the virtue of humility could essentially involve extended cognitive processes.

On the face of it, there seems no inherent reason why this shouldn't be so. Indeed, it might be that one could help someone both acquire and maintain this virtue by employing extended cognitive processes. In particular, there seems no inherent reason why technology cannot be employed to assist the cognitive processes that underlie the successful manifestation of virtue, and if that is right then so long as this technology becomes suitably cognitively integrated then we should be able to think of these cognitive processes as extended in the relevant sense.

For example, accounts of humility characteristically regard this virtue as demanding, *inter alia*, that one doesn't regularly overestimate one's importance. There would thus be a cognitive trait underlying this virtue that involved making judgements about one's importance that don't inaccurately overstate that importance.⁹ On the face of it, technology ought to be useful in assisting subjects in this regard. One could imagine, for example, a subject employing a device that is designed to pick-up on the use of certain phrases (in the subject's speech and writing, say) that are indicative of non-humble attitudes (arrogance, and so on). The device could then be programmed to remind the subject (in a suitably evocative manner) that one shouldn't overestimate one's importance. In this way, a general motivation to be humbler could be more effectively realised via the subject's continued employment of this technology.

Imagine now that this device is employed not as mere technology, but is instead cognitive integrated such that it becomes neuromedia. Perhaps, for example, that rather than consciously using this device, such that one is aware of it as an external instrument, one is fitted with it in some unobtrusive way, such that on a day-to-day basis one is not even aware of its presence. Relatedly, the way that it works, on suitable occasions, to generate thoughts that stimulate humility is such that the thoughts engineered by the device are not noticeably different from the various thoughts that pop into one's head during a normal day. We would thus have an example of neuromedia, in that the subject's use of this technology would be as seamlessly integrated into one's cognitive processes as one's entirely on-board cognitive processes, with a consequently similar associated phenomenology.

Interestingly, one would think that this device, *qua* neuromedia, would be even more effective at promoting humility than its non-extended counterpart. One would simply find oneself becoming humbler without even being aware of the role of the technology in this regard. In contrast, being aware of the technology and its role in one's beliefs could serve to undermine its effectiveness, since it naturally prompts the subject to reflect on the input from the technology *qua* the bearer of 'external' advice, rather than as 'internal' monitoring. For example, one might be inclined to downgrade the external input when it comes to guiding one's actions (i.e., what one will say or write next) precisely because it is external, preferring instead one's own internal judgement. In contrast, if the 'external' monitoring is not noticeably different from the 'internal' monitoring, such that the two are seamlessly merged, then the potential for the kind of intellectual distance required to downgrade the former over the latter simply doesn't arise.

Should one then conclude that it would be better for the technology that we employ to promote humility to be neuromedia where possible? One reason why we might pause here concerns the fact that virtues seem to demand a high level of cognitive ownership on the part of the subject. One does not acquire and maintain one's virtues passively, but rather actively, through conscious effort, emulation of the virtuous, reflection on one's performance to cultivate that virtue, and so on. The reason why this is relevant is that it suggests that there is a limit to the extent to which one can thoughtlessly manifest a virtue, which means that there is a sense in which one doesn't want one's employment of the cognitive processes that underlies the virtue to be *too* seamless. Accordingly, if it is in the nature of neuromedia that it is technology that is employed in a way such that one is unaware of employing it (at that moment at any rate), then one might think that this is a bar to it forming part of the cognitive basis for the manifestation of the virtue.

In order to see the import of this point, compare one's use of the unextended virtue-enhancing technology with one's use of the corresponding extended virtue-enhancing neuromedia. When one uses the unextended technology, one is consciously taking responsibility for one's development of one's virtue of humility. After all, one is actively employing the technology to enhance one's humility, and so there is a very real sense in which the technology is merely a tool for one's own cultivation of the virtue. Accordingly, one is able to take cognitive responsibility for the use of this technology in cultivating one's virtue. How does one's use of the neuromedia fare in this regard? The issue of cognitive responsibility is unclear when it comes to the use of neuromedia precisely because of how the technology is so seamlessly (and thus unreflectively) employed. Doesn't that suggest that one isn't now cultivating one's virtue at all, but is rather passively relying on the technology to engineer the relevant responses?

Another way of putting this point is in terms of where the causal attribution of responsibility naturally flows in each case. Where one is simply using technology as (unextended) technology in the usual way, then one's enhancement of the relevant virtue-associated behaviours naturally flows to the agent rather than to the technology, since the latter is merely being put into service by the former. But once the technology becomes neuromedia, then the relevant attribution of responsibility becomes much more muddled. In particular, if the responsibility for virtue-associated behaviours flows to the technology rather than to the agent, then it ceases to be straightforward that this behaviour is sufficiently creditable to the agent to qualify as the manifestation of virtue.

One could push back on this thought by emphasising that extended cognitive processes are simply parts of the extended cognitive subject, of her cognitive character.¹⁰ Accordingly, the idea that the virtue-associated behaviours are creditable to the technology rather than the subject is simply incoherent, as the technology is now part of the cognitive subject in the relevant sense. Whether this kind of push-back is credible depends on what counts as neuromedia. Where the subject actively incorporates neuromedia into her cognitive life with the expressed purpose of enhancing virtue, then it does seem very plausible to suppose that the unreflective use of the technology later in time is no less creditable to her agency for being employed without any conscious awareness of its use on the part of the subject. After all, part of what is involved in developing virtue is ensuring that the relevant behaviours and associated motivations become second nature, and hence one can hardly impugn one's choice to employ neuromedia in this fashion on the grounds that it leads to unreflective behaviour, as manifestations of virtue are ideally meant to be (in the moment at least) unreflective.

The more interesting case, however, is surely neuromedia that doesn't essentially involve any original conscious deliberation on the part of the subject. Although more controversial as an example of extended cognition than the type of neuromedia just described (whereby the subject does consciously choose to extend their cognitive processes in this way), it nonetheless is a plausible variety of *bona fide* extended cognition. Imagine, for example, a future world where everyone is automatically cognitively augmented from birth with certain kinds of neuromedia. Now consider an agent who is unaware of the nature of this cognitive augmentation, but nonetheless unthinkingly employs it in her everyday life. Insofar as we are inclined to treat neuromedia as a genuine extended cognitive process at all, then wouldn't we likewise be inclined to treat this subject as exhibiting extended cognition?¹¹ But if the subject has never consciously endorsed this use of technology, then in what sense would it be appropriate to treat the virtue-apt behaviours that result from the use of this technology as attributable to her agency (such that they could count as a genuine virtue)?

3. COGNITIVELY EXTENDED INTELLECTUAL HUMILITY

These issues about virtue and responsibility in the context of extended cognition become even more vexed once we turn our attentions from a general virtue like humility and focus instead on the potential role of extended cognitive processes with regard to the specifically intellectual virtue of intellectual humility. As noted above, since the intellectual virtues are themselves a kind of cognitive trait, then the issue is not merely whether the cognitive bases of the virtue can be extended, but whether the intellectual virtue itself can be cognitively extended. As we will see, the reasons why we should be cautious about thinking of a virtue like humility as even *involving* an extended cognitive process become even more pressing once we turn to the more ambitious idea of intellectual humility as actually *being* an extended cognitive process.

Indeed, building on our previous point about the virtues more generally and extended cognition, there seem to be good reasons for supposing that no intellectual virtue could be an extended cognitive process. If that's right, then at most the cultivation of intellectual humility could be in part aided via the employment of extended cognitive processes; it could never be the case that intellectual humility could be itself an extended cognitive process. For comparison, let's start with this latter, weaker, claim, in order to bring out what is problematic about the former, stronger, claim.

Just as we can imagine an agent making use of technology to aid her development of humility, even to the extent that the technology qualifies as neuromedia (though this is more controversial, as we saw above), so we can also imagine an agent making use of technology to aid her development of intellectual humility. As before, the exact manner in which this occurs will depend on what one holds intellectual humility to be, but in keeping with our remarks earlier we could treat it as relatively uncontroversial that intellectual humility at least demands that one doesn't regularly overestimate one's intellectual abilities and achievements.¹² Accordingly, we can imagine technology that encourages a healthy mindfulness in this regard, and in the process enables one to be more intellectually humble. Moreover, if this technology is developed along the right lines, then it may qualify as neuromedia, due to our seamless interactions with it and its cognitive integration with our other cognitive capacities. We thus have a case of the intellectual virtue of intellectual humility being supported by neuromedia, and thus incorporating the use of an extended cognitive process.

The interesting question, however, is whether we can move from this kind of claim to the stronger thesis that the intellectual virtue of intellectual humility can itself consist in an extended cognitive process, as opposed to merely incorporating such processes in its cultivation. We can see the implausibility of this suggestion by outlining what it might involve in particular cases. For example, suppose we fitted subjects with neuromedia from birth such that they exhibited behaviour associated with intellectual humility. No matter how cognitively integrated this neuromedia became with the subject's other cognitive processes, would there be any temptation to treat this behaviour as the manifestation of an intellectual virtue (rather than just the product of the technology)?

We can bring this point into sharper relief by imagining two groups of subjects, where the first group is trained up to be intellectually humble in the usual way (through regular emulation of intellectually humble exemplars, say), while the second group is simply fitted with the neuromedia to ensure that they end up exhibiting intellectually humble behaviour. Even if the behaviours ultimately exhibited by the two groups are identical (and in fact it is likely that the second group would exhibit the intellectually humble behaviours much more quickly), it only seems to be the former group who are genuinely manifesting intellectual humility. The crux of the matter is that intellectual virtues seem to be in part characterised by the essentially reflective manner of their acquisition and cultivation, such that even if they can be exhibited on particular occasions in an unreflective manner (as manifestations of one's second nature, whereby via habituation the exercise of virtue has become automatic), it nonetheless remains that no virtue, properly speaking, can be completely manifested

unreflectively. And yet that is exactly what the idea of neuromedia as a cognitively off-loading of intellectual virtue implies.¹³

One can imagine various kinds of critical push-back against this line. If cognitive abilities in general can become extended cognitive processes, then why not the specific kinds of cognitive abilities at issue when it comes to the intellectual virtues? (But then it has often been the case that the intellectual virtues have been explicitly contrasted with mere cognitive skills, just as the virtues in general are typically distinguished from mere skills).¹⁴ Or if one grants that there can be such a thing as extended cognitive abilities, then doesn't it follow that there can be an extended cognitive character? Accordingly, why not hold that the lines of cognitive responsibility do genuinely thread back to the extended cognitive agent, and hence that there is the relevant degree of cognitive responsibility involved on the part of the subject when neuromedia is in play that is needed for the manifestation of an intellectual virtue like intellectual humility? (But isn't the reality rather that a divide opens up between two kinds of cognitive responsibility: the lower-grade sort that is applicable to an extended cognitive character, and the higher-grade sort that is applicable to the intellectual virtues).

4. CONCLUDING REMARKS

This is not the place to settle these issues, but it is important to note that there are standing questions regarding the extent to which extended cognitive processes can be part of the manifestation of both the general virtue of humility and the specific intellectual virtue of intellectual humility. We have seen that there is some plausibility to the idea that the cognitive traits that underlie the general manifestation of virtue involved in humility could be extended. But we have also seen that there may be some limitations to the extent to which a virtue can rely on extended cognitive traits, due to the kind of cognitive responsibility involved in the manifestation of a genuine virtue. Moreover, this issue becomes more acute once we consider the question of whether intellectual humility could itself be an extended cognitive process.¹⁵

REFERENCES

- Adams, F., & Aizawa, K. (2008). *The Bounds of Cognition*, Blackwell, Oxford.
- Alfano, M., & Skorburg, G. (2016). 'The Embedded and Extended Character Hypotheses', *Routledge Handbook of Philosophy of the Social Mind*, (ed.) J. Kiverstein, 465-78, London: Routledge.
- Annas, J. (1995). 'Virtue as a Skill', *International Journal of Philosophical Studies* 3, 227-43.
- (2011). *Intelligent Virtue*, Oxford: Oxford University Press.
- Carter, J. A., Clark, A., Kallestrup, J., Palermos, S. O., & Pritchard, D. H. (2014). 'Varieties of Externalism', *Philosophical Issues* 24, 63-109.
- (Eds.), (2018). *Socially Extended Epistemology*, Oxford: Oxford University 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, 7-19.
- Giere, R., & Moffat, B. (2003). 'Distributed Cognition: Where the Cognitive and the Social Merge', *Social Studies of Science* 33, 1-10.
- Gallagher, S. (2013). 'The Socially Extended Mind', *Cognitive Systems Research* 25/26, 4-12.
- 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.
- 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.
- Priest, M. (2017). 'Intellectual Humility: An Interpersonal Theory', *Ergo* 4 (DOI: <http://dx.doi.org/10.3998/ergo.12405314.0004.016>).
- Pritchard, D. H. (2010). 'Cognitive Ability and the Extended Cognition Thesis', *Synthese* 175, 133-51.
- (2014). 'Truth as the Fundamental Epistemic Good', *The Ethics of Belief: Individual and Social*, (eds.) J. Matheson & R. Vitz, 112-29, Oxford: Oxford University Press.
- (2016). 'Epistemic Axiology', *Epistemic Reasons, Epistemic Norms, and Epistemic Goals*, (eds.) M. Grajner & P. Schmechtig, 407-22, Berlin: DeGruyter.
- (2018a). 'Extended Knowledge', *Extended Epistemology*, (eds.) J. A. Carter, A. Clark, J. Kallestrup, S. O. Palermos & D. H. Pritchard, 90-104, Oxford: Oxford University Press.
- (2018b). 'Extended Virtue Epistemology', *Inquiry* 61, 632-47.
- (2018c). 'Intellectual Humility and the Epistemology of Disagreement', *Synthese* [Online First, DOI: <https://doi.org/10.1007/s11229-018-02024-5>].
- (2018d). 'Neuromedia and the Epistemology of Education', *Metaphilosophy* 49, 328-49.
- (2019). 'Disagreement, Intellectual Humility, and Reflection', *Thinking about Oneself: The Place and Value of Reflection in Philosophy and Psychology*, (ed.) W. Da Silva Filho, chapter 5, Dordrecht, Holland: Springer.
- Rupert, R. D. (2009). *Cognitive Systems and the Extended Mind*, Oxford: Oxford University Press.
- Stichter, M. (2018). *The Skillfulness of Virtue: Improving our Moral and Epistemic Lives*, Cambridge: Cambridge University Press.
- Tanesini, A. (2016). 'Intellectual Humility as Attitude', *Philosophy and Phenomenological Research* (Online First, DOI: 10.1111/phpr.12326).

- Theiner, G., Allen, C., & Goldstone, R. (2010). 'Recognizing Group Cognition', *Cognitive Systems Research* 11, 378-95.
- Whitcomb, D., Battaly, H., Baehr, J., & Howard-Snyder, D. (2017). 'Intellectual Humility: Owning our Limitations', *Philosophy and Phenomenological Research* 94, 509-39.
- 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 Note that for the purposes of this article I will be setting to one side a particular kind of extended cognition—sometimes called *distributed* or *socially extended cognition*—where the cognitive extension involves other agents. For some useful discussions of distributed cognition, see Giere & Moffat (2003), Theiner *et al* (2010), and Gallagher (2013). For specific discussion of the epistemological ramifications of distributed cognition, see Palermos & Pritchard (2016) and Carter *et al* (2018).

² Note that Clark & Chalmers (1998) are less interested in arguing for extended cognition than for the more specific thesis of the *extended mind*—*viz.*, roughly, that minds can extend beyond the skull of the subject. Whether the extended mind thesis is a more demanding thesis than the extended cognition thesis—as Clark & Chalmers (1998) clearly believe—depends on further factors, such as whether one treats the cognitive as the mark of the mental. In any case, our concern here is with the extended cognition thesis rather than the extended mind thesis. For Clark’s more considered take on extended cognition (/mind), see Clark (2008).

³ Note that the ‘neuromedia’ terminology is not my invention, but since this terminology is used in varied ways it is important that it is understood along the specific lines set out here. For a philosophical discussion of neuromedia that is relevant to our current concerns (but which uses this terminology in a slightly different manner), see Lynch (2014; 2016).

⁴ Note that it might be important to the technology that the subject *can* tell the difference, if only in principle, perhaps for legal reasons. But it would still be significant that there could exist technology of this kind where ordinarily the subject could not tell the difference.

⁵ Interestingly, one way in which neuromedia might be developed could be as devices that are fitted beneath the skin and skull of the subject. Indeed, if one is aiming for the technology to be as seamlessly integrated into one’s cognitive character as much as possible, such that one isn’t even normally aware of it as technology, then it would make sense to have it as hidden from view as much as possible. If neuromedia is developed in this way, however, then it is in a certain sense an ‘internal’ form of extended cognition (even though it is still ‘external’ in the manner that is relevant to extended cognition, as it is external the biological cognitive processes that the subject has under her skin and skull).

⁶ For further discussion of cognitive integration in the context of extended cognition, see Palermos (2014*a*; 2014*b*).

⁷ For two prominent critiques of extended cognition—and also the related extended mind thesis (see endnote 2)—see Adams & Aizawa (2008) and Rupert (2009).

⁸ That’s the standard way of thinking about the intellectual good at any rate—see, for example, Zagzebski’s (1996) influential account of the intellectual virtues in this regard—though some might wish to substitute a more elevated epistemic standing like knowledge, understanding or wisdom. (For what it is worth, I think that this would be a mistake—see Pritchard (2014; 2016)).

⁹ I leave it as an open question whether this cognitive trait would be compatible with one making *inaccurate* judgements that regularly *underestimate* one’s importance—this would depend on further features of one’s account of humility.

¹⁰ For a development of the idea of an extended cognitive character, see Alfano & Skorburg (2016).

¹¹ I discuss a case of this kind, whereby the extended cognitive process is never consciously endorsed by the subject, but where it nonetheless seems to count as a *bona fide* instance of extended cognition, in Pritchard (2010).

¹² There are a range of different accounts of intellectual humility in the contemporary literature. For some key recent discussions in this regard, see Tanesini (2016), Whitcomb *et al* (2017), and Priest (2017). For my own take on these issues, see Pritchard (2018*c*; 2019).

¹³ I develop this idea in Pritchard (2018*a*; 2018*b*; 2018*d*).

¹⁴ For some recent treatments of virtues, including intellectual virtues, as being a kind of skill—which is, of course, compatible with the idea that intellectual virtues might be more demanding cognitive traits than other kinds of (mere) cognitive skill—see Annas (1995; 2013) and Stichter (2018).

¹⁵ I am grateful to Mark Alfano for detailed comments on an earlier version of this paper.