

# Computer Ethics - Philosophical Enquiry (CEPE) Proceedings

---

Volume 2019 *CEPE 2019: Risk & Cybersecurity*

Article 4

---

5-30-2019

## Hacking the Extended Mind: The Security Implications of the New Metaphysics

Robin L. Zebrowski  
*Beloit College*

Follow this and additional works at: [https://digitalcommons.odu.edu/cepe\\_proceedings](https://digitalcommons.odu.edu/cepe_proceedings)



Part of the [Information Security Commons](#), and the [Metaphysics Commons](#)

---

### Custom Citation

Zebrowski, R. L. (2019). Hacking the extended mind: The security implications of the new metaphysics. In D. Wittkower (Ed.), *2019 Computer Ethics - Philosophical Enquiry (CEPE) Proceedings*, (12 pp.). doi: 10.25884/sy6p-ch18 Retrieved from [https://digitalcommons.odu.edu/cepe\\_proceedings/vol2019/iss1/4](https://digitalcommons.odu.edu/cepe_proceedings/vol2019/iss1/4)

This Paper is brought to you for free and open access by ODU Digital Commons. It has been accepted for inclusion in Computer Ethics - Philosophical Enquiry (CEPE) Proceedings by an authorized editor of ODU Digital Commons. For more information, please contact [digitalcommons@odu.edu](mailto:digitalcommons@odu.edu).

# Hacking the Extended Mind: The Security Implications of the New Metaphysics<sup>1</sup>

Robin L. Zebrowski  
Beloit College

## Abstract

Computer security expert Paul Syverson has argued that there is a computer security equivalent of gaslighting: where a clever adversary could convince some system that some component that is not really a part of the system is in fact a part of the system. If non-biological items from our environments (or even pieces of our environments themselves) can be part of our minds (the standard Extended Mind hypothesis, EM), they are therefore part of our selves, and therefore subject to Syverson's worry about boundary in a way that has not been explored before. If some version of EM holds, then what were once security concerns surrounding various systems or devices become those same concerns but writ large for our cognitive processes, the core of mind and thought. Philosophers and critics have long been worried that if EM is true, selfhood bleeds out of the nicely-contained package of skin and skull. Criticism has been offered that argues if we allow some of the environment to count as genuinely part of our minds, it seems to threaten the existence of our phenomenal feeling of ownership over our own bodies, and the worry persists that selfhood itself becomes threatened with dissolution. Yet if we agree with Clark and understanding selfhood as fundamentally dispersed and therefore only imperiled by the more traditional view of some core, persistent notion of self, the security worry remains. No matter how we adjust our foundational metaphysics to account for the extended mind hypothesis, we are stuck with the security concerns, and we must start to catalog and account for the challenges before we see actual adversarial attacks on the actual devices constituting the minds of real people. Thinking about the very nature of the mind turns out to provide insight into how we think about security, and vice versa.

*Keywords: extended mind, extended self, security*

One random day in 2019, at 1.50am, I received a notification on my phone, *from* my phone. It said, simply, "You have a new memory- June 16, 2016." The phone wasn't just reporting on some organizational criteria set previously, and it wasn't some special app designed for nostalgia; it was the stock "photos" app, inventing a new gallery of three-year old photos, presumably to remind the user of an event that some algorithm decided reached a threshold of engagement. But the experience served its purpose: the images showed a fun day that had been long forgotten, and so the phone was not wrong in its

---

<sup>1</sup> The central theses of this paper were suggested to me by Paul Syverson, and its content is the result of extensive discussions with and comments from him. I owe him a great debt in helping me refine these ideas. This version of this paper is submitted "as presented" at CEPE 2019 in Norfolk, Virginia at Old Dominion University, and is meant to be a snapshot of the conference presentation more than a complete and robust argument. A more detailed version is forthcoming.

declaration of a “new memory.” I looked at the phone, recalled a pleasant visit to a beloved ice cream shop in NJ with my then-toddler, and recognized with trepidation that it was, in fact, a new memory. Was the word “memory” a misnomer applied in this way? And if not, was the memory generated by the phone, or stored by it, the twenty-first century augmented version of work usually done by a brain?

In the more than twenty years since Clark and Chalmers published “The Extended Mind,” there have been dozens of volumes written exploring the veracity of the claims they made as well as the potential implications for a number of branches of philosophy. Here, I present an area of impact (security) that has been under-explored, with a novel approach that I hope will lead to new ways of thinking about both the mind/self and new ways of approaching security in relation to certain devices and aspects of our environment that might rightly be said to play a constitutive role in our selfhood. This paper will describe the basic claims of the extended mind hypothesis, discuss the relationship between mind/self, both ontologically and in the extended mind literature, and then introduce the notion of gaslighting in relation to systems (broadly) with a self-concept (broadly). I’ll then argue that there are multiple metaphysical pictures that are consistent with the extended mind hypothesis, and these security worries remain regardless of which version holds.

## **Extended Mind Hypothesis**

Clark and Chalmers (1998) offer an illustration in order to introduce the possibility that we (in the West, anyway) have been largely mistaken about the boundaries of the mind. They propose thinking about three different ways we might interact with a Tetris-like game, regarding how we manipulate the pieces to test for spatial fit (7). In the first case, the subject uses mental rotation to determine fit. In the second case, the subject has a choice to either use mental rotation or to use a rotate button that displays on the screen. They correctly speculate that using the physical rotate button might confer some speed advantages over mental rotation. In the third case, we are to imagine a future wherein the subject has a neural implant which is as fast as the physical rotation option in the second case, but which the subject has the option to use or not use in any given scenario. They then ask how much cognition is present in each of these cases. The argument they offer equates the third case (neural implant) with the first case (mental rotation), and the third case with the second case (physical button + screen), and concludes that if we are willing to allow the third case to count as cognitive (it is internal, neural, computationally similar) then it would seem the second case is equally cognitive (it just happens to loop in the hand, the keyboard, and the screen in a slightly different way). If you are willing to allow this possible cognitive equivalence, then the authors have successfully challenged the notion that brain and mind are the same thing, or even that body and mind are the same. Instead, they claim, properly understood, certain tools in our environments are also genuinely cognitive, not merely causally but constitutively, meaning that one’s mind is not contained to one’s body in some cases.

The authors offer another scenario to illustrate the general argument by introducing the fictional characters of Otto and Inga (ibid, 12). Briefly, Inga is neurotypical, and Otto has Alzheimer’s. We might imagine that Inga wishes to visit the

museum, thinks about it for a moment to remember where it is, recalls the museum is on 54<sup>th</sup> Street, and then goes there. As a result of his memory difficulties, Otto carries a notebook with him everywhere he goes, wherein he writes down important information he needs to remember. He, too, wants to visit the museum, consults his notebook, learns the museum is on 54<sup>th</sup> Street, and off he goes. The important move here comes via the notion of belief: Inga believes the museum is on 54<sup>th</sup> while she's on her way there, but it would be wrong to claim that she doesn't have such a belief even when she isn't thinking about it. (She has a non-occurrent belief before she calls up its location in her memory.) The notebook, then, is claimed to serve the same function for Otto as Inga's biological memory serves for her. It should follow, then, that Otto, too has a non-occurrent belief that the museum is on 54<sup>th</sup>, and like Inga's belief, his becomes occurrent when he looks it up and sets off.

The tricky part, then, is the claim that the belief lives outside of Otto's head, in the notebook. Again, volumes have been written on this claim (Adams and Aizawa 2008, Menary 2006, Clark 2008, Estany and Sturm, 2014, Robbins and Aydede 2009, Schantz, 2004), so it won't be heavily defended here. I will simply assume the extended mind hypothesis is plausible. Clark and Chalmers refer to this idea as the Parity Principle, and it is understood to be a rule of thumb rather than a law of nature: if something would be called cognitive, were it done in the head, then it should also be called cognitive when it uses tools external to the body. If I use a pen and paper to compute a long division problem, but you are a math whiz who can do the entire computation unaided ("in your head"), then we ought to consider my brain/body plus the pen and paper cognitively equivalent to your brain/body. There are a number of reasons we might care about the veracity of the claim. Clark (2003) argues, "We exist, as the thinking things we are, only thanks to a baffling dance of brains, bodies, and cultural and technological scaffolding. Understanding this evolutionarily novel arrangement is crucial for our science, our morals, and our self-image both as persons and as a species" (11). The sense of self that he refers to here and elsewhere (5) is part of what is at stake in my arguments here.

## **Mind/Self**

Unless you are a very strict Cartesian, you probably don't equate your mind and your self. While the notion of the extended mind has been heavily explored, the idea of the extended self has been somewhat less so (at least in these terms)<sup>2</sup>. Where mind and self are collapsed, they are mostly collapsed for historical reasons (see, again, Descartes, who really screwed us.) I hope to engage in a more extended argument for the extended self in a future version of this paper, but for now I look to the original Clark and Chalmers paper again for validation that the extended mind likely also entails an extended self: "Does the extended mind imply an extended self? It seems so. Most of us already accept that the self outstrips the boundaries of consciousness; my dispositional beliefs, for example, constitute in some deep sense part of who I am. If so, then these boundaries may also fall beyond the skin. The information in Otto's

---

<sup>2</sup> For one important argument for the extended self, see Malafouris (2008)

notebook, for example, is a central part of his identity as a cognitive agent. What this comes to is that Otto *himself* is best regarded as an extended system, a coupling of biological organism and external resources”(18). It seems that Clark and Chalmers think the extended mind is primary, and selfhood likely follows. Perhaps there is an argument to be made that the notion of the extended self does not rely on the notion of the extended mind at all, and Otto’s notebook might still be a candidate case. While I will not argue this here, I find it plausible.

## Functionalism

Since rising to prominence in the late 1950s/early 1960s, functionalism has remained one of the most popular views of mind in philosophy. Importantly, the relationship of the extended mind to functionalist views is somewhat messy and hotly debated. The basic claim of the functionalist view is that the nature of mind rests on the causal relationships between various mental states, inputs, and outputs, as opposed to something like reductive materialism, where the mental states are instead considered identical to some specific physical state (this is the view that would argue, for example, that a neural firing is identical to seeing the color blue. Functionalism, instead, would talk about some set of causal relations, usually starting with some stimulus in the world acting on some blue-detecting system, producing some outcome such as “says, ‘this is blue’”.)

This paper does not intend to defend or attack functionalism, but mentions it because of an ongoing debate about whether functionalism itself logically entails the extended mind/self. This debate is only relevant insofar as the functionalist view remains popular, and many who would deny the extended mind hypothesis would fully endorse functionalism, and it isn’t clear whether this is a consistent view to hold or not. The debate is again beyond the scope of this paper, but the range of ways of framing the problem is at least curious, and perhaps this paper’s argument rests on one of these claims in ways I have not fully considered yet. Sprevak (2009) argues that functionalism does entail the extended mind, and therefore concludes that functionalism itself is proven wrong as a result of such an absurd conclusion. Wheeler (2010) instead agrees that the extended mind is entailed by functionalism, but then says that the result is how we know both of these views must be correct, since functionalism is correct. Adams and Aizawa (2001) have argued that functionalism does not entail the extended mind, but they also do not think that it precludes it (they just haven’t yet seen evidence for the extended mind). Rupert (2004), too, argues that functionalism does not entail the extended mind, but also thinks that it doesn’t preclude it. Rounding out all possible views, Miyazono (2015) argues that functionalism entails that the extended mind must be false.

Even in its earliest explorations, the extended mind was lurking in discussions of functionalism. In his famous paper arguing against Functionalism, Block acknowledges, “if we are ‘full power’ Turing Machines, the environment must constitute part of the tape” (213). It remains somewhat surprising that forty years after Block wrote this, the relationship between the extended mind and functionalism remains so complicated.

Again, I mention functionalism and the debate about its relationship to the extended mind hypothesis largely to dissipate any immediate pushback on the idea

from the standpoint of functionalism. As we can see, the debate is ongoing and consensus seems far from settled on the matter. I believe (or hope) that this means I do not need to come down in any settled manner for or against any version of functionalism in order to make this argument about the extended self and its complicated relationship to security.

## **Systems Security and Gaslighting**

If you have a system that has some sort of self-representation, then an adversary might attempt to convince the system that some component that is not really a part of the system is in fact part of the system. The system then might trust and protect the adversarial piece as a part of itself. The system may then end up trusting something malicious or protecting something foreign that is not a legitimate part of the system in question. Conversely, an adversary might convince the system that some component of itself is not part of the system. The system might then ignore inputs from that component or indicators about the state of that component as not self-relevant. In the case of the extended self, we are such systems, and our external devices are potentially targets in a way not previously appreciated. Successful attacks then amount to a form of gaslighting, manipulating someone to intentionally cause them to doubt their own perceptions or memory.

If you create confusion about the boundaries of the self of the system, a number of possibilities follow. You can make it ignore some parts of itself. You can make it treat some part of itself as not-itself. You can make it treat some part of not-itself as itself. You can make it treat some part of itself as another. In each case, the adversary is then well-positioned to leverage the confusion.

Cooley (1902) famously remarked that, “I am not what I think I am; and I am not what you think I am; I am what I think you think I am.” Thus, to some extent, your self-concept already encompasses some adversaries; they help define who and what you understand yourself to be. A mouse, for example, has a very different representation of itself than I would in an environment where it perceives a cat perceiving it. Its survival is entwined with its sense of cats as adversaries in a way mine is not.

If the extended self is plausible (which it seems to be), and this framework for understanding adversarial relations to a system using a technique like gaslighting is plausible (which it seems to be), then this same framework applies easily to our external suite of tools when they meet the conditions to count as part of ourselves. For example, women often rely heavily on period tracker apps on smartphones, offloading vital memory and prediction information to such apps and trusting the information they supply in return. In a May 2019 article, the Guardian reported that one such app was created and supported by opponents of abortion and had been encouraging women to forego hormonal birth control in favor of trusting the app and its predictions about fertility in their quests to avoid (or achieve) pregnancy (Glenza, 2019). They go on to report, “two of the app’s medical advisors are not licensed to practice in the US and are also closely tied to a Catholic university in Santiago, Chile, where access to abortion remains severely restricted.” The argument that the smart phone may qualify as an extension of myself is established, if controversial, and certain apps on such a phone would likely

qualify even more clearly, particularly in a case like this, where this is cognitive work related to self-preservation that is often offloaded onto external tools. Eva Galperin, a security researcher at the Electronic Frontier Foundation, in discussing the threats posed by online stalkers, noted that, “full access to someone’s phone is essentially full access to someone’s mind” (Greenberg, 2019).

We can imagine a number of gaslighting-style errors that we make on a bodily level that might help support the claims I want to make about adversarial worries with regard to the extended mind. Certain kinds of visual illusions would seem to fall under this umbrella, as for example when some aspect of our visual system becomes fatigued and it produces an effect that fools us into imagining we see movement where there is none, or in the Hermann Grid illusion, where darkness appears where only white space exists. These effects happen as a normal by-product of our visual systems functioning correctly; they aren’t errors, but instead could be described as our bodies tricking themselves as a result of their otherwise-normal behavior. This would be a gaslighting-style error insofar as anyone with a basic understanding of these sorts of visual processes could (hypothetically) leverage that exploit against us, having our eyes report that we are seeing something that we aren’t actually seeing. Another example might involve clever pickpocketing techniques. Paul Syverson tells a story about how he witnessed an expert pickpocket (Apollo Robbins) describe his trade in front of an audience. When the pickpocket introduced himself to the target, he grabbed his shoulder and shook his hand, and otherwise touched him a bit more forcefully than is normally expected (but not so far out of the bounds of normal as to stand out as inappropriate to the target). The pickpocket was then able to remove the target’s watch in front of a live audience without the target noticing. In effect, what the pickpocket did was fatigue the touch sensors of the target’s body, such that a very light touch didn’t register and rise to conscious awareness, because the target had felt such strong pressure during the greeting. This is a kind of attack on the self-model of the target insofar as it’s training him to alter that model via the relative significance attached to those signals. This is obviously exploitable, but it is also a gaslighting-style error on a bodily level, where a person is not experiencing the touch on their body as touch on their body. Therefore, it’s easy to see how this kind of error can happen on the bodily-level. But what about the tool-level?

As long as we restrict the tools here to those tools relevant to your self-concept in an extended system, we can imagine several concrete ways this might work. In addition to the app mentioned above, we can also imagine messaging systems, when they are well-integrated into your daily life and existence in a way that qualifies under the extended mind/self conditionals, that may exploit this access in a similar way. A message may arrive purporting to be from some trusted source, but it is not from such a source. Such messages (and the systems that enable them, again only when those systems are well-integrated into your self-concept) may qualify as this kind of exploit. The argument here is not that there is something novel in this kind of exploit, popular for many decades, but that when those devices and systems are part of the extended self,

the exploit serves a more serious kind of attack *on the person* than previously recognized.<sup>3</sup>

We might also think of the way some machine learning algorithms are being automatically endorsed by people in certain situations such that they come to rely on the output of such algorithms in their everyday lives or work without considering double-checking on that output.<sup>4</sup> But of course, the ways those algorithms can already be attacked and exploited is fairly well-known. We saw those attacks appear as altered pixels in image-recognition software that caused algorithms to mis-characterize images that look (to humans) like one thing, but to the software like something else. This can go beyond purely digital manipulation, working in the real world in various ways, too. For example, it was reported that with a few simple stickers, a stop sign was read by such an algorithm as a speed limit sign (Hutson, 2018), which can have all sorts of effects on autonomous cars (and our relationship to our cars is already understood as one of extension, broadly understood). Similar attacks have been carried out using vocal samples instead of visual ones, with researchers at Google reporting they were able to add inaudible elements “that sounded to humans like ‘without the data set the article is useless,’ but that an AI transcribed as ‘OK Google, browse to evil.com’” (ibid). These sorts of adversarial attacks are replicable and fairly-well understood by security researchers, but we haven’t considered the ways they might take on additional weight as they are deployed against personal devices or systems that comprise part of our extended selves.

Since most of the literature that has dealt with the extended mind since the original publication of Clark and Chalmers (1998) has focused on the argument as it relates to mind, it is difficult to find critique of this extended view as it relates to both selfhood and security. However, Sterelny (2004) comes close, worrying about Otto (who he refers to as “Mr T”) and his notebook being potentially subject to what he calls “epistemic sabotage” (246). His argument largely rests on the notion that since Otto’s notebook is subject to manipulation and sabotage in a way Inga’s biological memory apparently isn’t, it cannot be equivalent enough in a way that matters to count as part of his extended mind. Sterelny worries about “thought insertion,” by way of someone writing in the notebook that Otto will later automatically endorse, as well as Otto having his beliefs stolen. He points out that “this set of problem[s] simply does not arise for such of Mr T’s information that he still encodes internally” (246). Clark uses this opportunity to double down on the arguments for extension, saying, “The point about vulnerability to malicious manipulation is well taken. Many forms of perceptual input are indeed subject, for that very reason, to much vetting and double-checking. I do not think, however, that we treat all our perceptual inputs in this highly cautious way. Moreover, as soon as we do not do so, the issue about extended cognitive systems seems to open up...” (60, 2006). In other words, everyone agrees that devices or systems that are not accessed introspectively as are our biological beliefs, for example,

---

<sup>3</sup> There is more to be explored here on the legal ramifications of these claims. See, for example, Stephen B. Wicker, “Smartphones, Contents of the Mind, and the Fifth Amendment,” (April 2018), *Communications of the ACM* 16(4). A deeper discussion of this is forthcoming.

<sup>4</sup> In the original 1998 paper, Clark and Chalmers argue that four criteria must be met for some external thing to count as part of the extended mind (17). One of these is that when the information is available, it is automatically endorsed by the subject.



are already subject to additional scrutiny. But Clark points out that not all externally-accessed perceptual and biological systems are subject to such scrutiny. The aforementioned visual illusions, for example, are taken in by our bodies as immediate experiences of perception, and it takes additional cognitive analysis to begin to make sense of the ways our immediate experiences may not be trustworthy. The same would surely be true of Otto's notebook, if he finds himself following its directions and recognizes that he is not at the museum but instead somewhere obviously different. We can understand a difference between our trust of something biological and some external artefact, but there are no cases for which we are always correct in that trust. As a result, it isn't clear that our privileging of biology over artefact is somehow more accurate or natural or correct, particularly in making sense of selfhood. Clark continues, "For the reason we are vulnerable in just those kinds of cases is, I would argue, because we are relying on an ecologically sound strategy of treating the external scene as a stable, reliable substitute for internally-stored memory traces. In short, our brains have decided (if you will allow such loose talk for a moment) that on a day to day basis the chances of these kinds of espionage are sufficiently low that they may be traded against the efficiency gains of treating the perception involving loop as if it were an inner, relatively noise-free channel, thus allowing them to use the world as 'external memory'" (60, 2008). Rather than using Sterelny's worry about the difference in reference between the internal and external, Clark takes this to be further evidence for his claims with regard to the external mind (and which I am arguing also holds for the external self). He concludes, "But what emerges is not so much an argument against the extended mind as a way of further justifying our claim that in some contexts signals routed via perceptual systems are treated in the way more typical of internal channels (and vice versa, in the case of standard thought insertion). To decide, in any given case, whether the channel is acting more like one of perception or more like one of internal information flow, look (in part) to the larger functional economy of defenses against deception. The lower the defenses, the closer we approximate to an internal flow" (61). Therefore, even Otto's notebook being tampered with isn't a clear reason to believe it is deeply different than biological memory.

## **Metaphysics is a Harsh Mistress**

This entire project rests on a few assumptions about the nature of embodiment and mind, namely, that the nature of the mind is not confined to the brain or to some ethereal Cartesian other-place of pure symbolic representation, but is instead something inherently embodied, embedded in a physical, social, and cultural world, and extended into some pieces of those environments in various ways under certain conditions. This holds also for the self, as I've argued. I find myself fighting a traditional Western understanding of mind and self that feels implicit in almost all work and is often stated outright in traditional scholarship about mind and self. However, I recognize that this assumption about a solitary, core self, once made explicit, often lurks in contemporary work, unargued for. In her discussion of virtual reality and embodiment (1999), Katherine Hayles offers a challenge to this underlying metaphysics that points out that people only feel threatened by something like the extended mind if their

underlying metaphysical worldview insists on mind or self being that sort of unchanging, isolated core. She says, "As long as the human subject is envisioned as an autonomous self with unambiguous boundaries, the human-computer interface can only be parsed as a division between the solidity of real life on one side and the illusion of virtual reality on the other, thus obscuring the far-reaching changes initiated by the development of virtual technologies. Only if one things of the subject as an autonomous self independent of the environment is one likely to experience the panic... This view of the self authorized the fear that if the boundaries are breached at all, there will be nothing to stop the self's complete dissolution. By contrast, when the human is seen as part of a distributed system, the full expression of human capability can be seen precisely to depend on the splice rather than being imperiled by it" (290). In other words, the entire fear of extension in relation to selfhood is completely dissolved if our starting assumption is one of distributed systems instead. But since we in the West rarely begin with that view, the extended self picture is one which receives a lot of pushback and fear.<sup>5</sup> It is our naïve metaphysics of individualism which drives much of the worry with this worldview. As John Dewey wisely remarked, it was as if "they took the structure of discourse for the structure of things" (170). This mistake drives many of our confusions.

There are, of course, many metaphysical pictures for which this sort of extended self is not such a challenge. Louise Barrett, for example, points out that "...we can put it another way: if we think of cognition as an active process, and "mind" as something animals do rather than something they "have," then questions about whether "minds" are things inside the head, or things that can exist outside them, don't really make much sense" (199). There is extensive work in process metaphysics that would support the ideas of a distributed or extended self-system. A full survey of such views is beyond the scope of this paper, but it is well-worth noting that these views are robust and historically strong, although generally under-explored in contemporary philosophy and psychology.<sup>6</sup> If we start from a process view, it isn't clear that the concerns about selfhood or mind and its lack of containment within a biological body and brain even arise, but, importantly, the security concerns still do.

## Security Concerns Remaining

While it is unlikely that the metaphysics of selfhood has changed as our devices have evolved, our understanding of that relationship must. If we are extended selves, genuinely comprised of both the biology and the relevant pieces of the environment, we likely have been this kind of thing for a long historical time. (Indeed, there is a fascinating literature emerging that traces this view of extended selfhood backwards via material culture, at least several millennia.)<sup>7</sup> If the argument here is correct, that these security concerns are not new, then we ought to be able to engage with some of that

---

<sup>5</sup> Non-Western views more often start with Hayles's more distributed understanding of self, or emphasize the idea that self is an illusion to begin with, and so we can use the label as needed for the sake of convenience.

<sup>6</sup> While you might trace process metaphysics in the West to Heraclitus, for the most robust versions, see Whitehead, for a historical approach, and Bickhard, for a robust contemporary one.

<sup>7</sup> See Malafouris 1997; 2008; 2013; Malafouris and Renfrew (eds.) 2010

archaeological literature, and read into it ways that security could have been or was a worry, even if it was unacknowledged. Iphones and Google did not create these security issues; they merely offer us a unique bit of insight into our own ways of conceiving ourselves, while highlighting for us ways that we must be vigilant going forward, as these particular devices at this particular point in history take up the role of partly enacting distributed pieces of our selfhood. Metaphysics remains a harsh mistress, but perhaps this time we can be prepared to defend whatever parts of ourselves we determine are worth becoming.

## References

- Adams, F. and Aizawa, K. (2001). "The Bounds of Cognition," *Philosophical Psychology* 14 (1): 43-64
- Adams, F. and Aizawa, K. (2008). *The Bounds of Cognition*, Oxford: Blackwell Publishing.
- Barrett, L. (1999). *Beyond the Brain: How Body and Environment Shape Human and Animal Minds*. Princeton: Princeton University Press.
- Bickhard, M. (2012) "A process ontology for persons and their development," *New Ideas in Psychology* 30: 107–119.
- Block, N. (1991). "Troubles with Functionalism" in *The Nature of Mind*, Rosenthal, D. (ed.) Oxford: Oxford University Press, 211-228.
- Clark, A. and Chalmers, D. (1998) "The Extended Mind," *Analysis* 58 (1): 7-19.
- Clark, A. (2006) "Memento's Revenge: The Extended Mind, Extended," in *The Extended Mind*, Menary, R. (ed.), Cambridge: MIT Press, 43-66.
- Clark, A. (2003) *Natural Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. Oxford: Oxford University Press.
- Clark, A. (2008) *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.
- Cooley, C. H. (1902) *Human Nature and the Social Order*. New York: Scribner's.
- Dewey, J. (1925/1958) *Experience and Nature*. Dover.
- Drayson, Z. (2010). "Extended Cognition and the Metaphysics of Mind," *Cognitive Systems Research* (11): 367-377

- Estany, A., and Sturm, T. (2014). Extended cognition: New philosophical perspectives. Special Issue of *Philosophical Psychology* (27): 1
- Glenza, J. (2019). "Revealed: Women's Fertility App is Funded by Anti-Abortion Campaigners." <https://www.theguardian.com/world/2019/may/30/revealed-womens-fertility-app-is-funded-by-anti-abortion-campaigners>
- Greenberg, A. (April 3, 2019). "Hacker Eva Galperin Has a Plan to Eradicate Stalkerware." <https://www.wired.com/story/eva-galperin-stalkerware-kaspersky-antivirus/>
- Haugeland, J. (1985). *Artificial Intelligence: The Very Idea*. Cambridge: MIT Press.
- Hutson, M. (July 20, 2018). "Hackers Easily Fool Artificial Intelligences," *Science* (361) 6399: 215.
- Malafouris, L. (2008). "Between Brains, Bodies and Things: Tectonoetic Awareness and the Extended Self," *Philosophical Transactions of the Royal Society B* (363): 1993-2002.
- Malafouris, L. and Renfrew, C. (eds.) (2010). *The Cognitive Life of Things: Recasting the Boundaries of the Mind*. Cambridge: The McDonald Institute for Archaeological Research
- Malafouris, L. (2013). *How Things Shape the Mind*. Cambridge: MIT Press.
- Menary, R. (ed.) (2010). *The Extended Mind*. Cambridge: MIT Press.
- Miyazono, K. (2015). "Does Functionalism Entail Extended Mind?" *Synthese* DOI: 10.1007/s11229-015-0971-2
- Robbins, P. and Aydede, M. (eds.) (2009) *The Cambridge Handbook of Situated Cognition*. New York: Cambridge University Press.
- Rupert, R. D. (2004). "Challenges to the Hypothesis of Extended Cognition," *The Journal of Philosophy* (101) 8: 389-428.
- Schantz, R. (ed.) (2004) *The Externalist Challenge: Current Issues in Theoretical Philosophy* Volume 2. Berlin: Walter DeGruyter
- Sprevak, M. (2009). "Extended Cognition and Functionalism," *The Journal of Philosophy* (106): 503-527
- Sterleny, K. (2004). "Externalism, Epistemic Artefacts and the Extended Mind," in *The Externalist Challenge*, Schantz, R. (ed.) De Gruyter: 239-254.

Wheeler, M. (2010). "In Defense of Extended Functionalism," in *The Extended Mind*, Menary, R. (ed.) Cambridge: MIT Press, 245-270.

Whitehead, A. N. (1929). *Process and Reality: An Essay in Cosmology*, New York: Macmillan

Wicker, S. B. (April 2018). "Smartphones, Contents of the Mind, and the Fifth Amendment," in *Communications of the ACM (61)4*: 28-31.