

Data, Metadata, Mental Data? Privacy and the Extended Mind

S. Orestis Palermos
University of Ioannina

Abstract. It has been recently suggested that if the Extended Mind thesis is true, mental privacy might be under serious threat. In this paper, I look into the details of this claim and propose that one way of dealing with this emerging threat requires that data ontology be enriched with an additional kind of data—*viz.*, mental data. I explore how mental data relates to both data and metadata and suggest that, arguably, and by contrast with these existing categories of informational content, mental data should not be merely legally protected. Rather, if we value mental privacy as we know it, technological measures should be employed to ensure that one's mental data are *practically*—not just legally—impossible for others to obtain.

Keywords: Extended Mind; Privacy; Data; Metadata

1. Introduction

Eavesdropping, wiretapping, snooping into one's personal belongings are disturbing breaches of privacy. But how do they compare with the possibility of having another looking into your very mind; peeking into your detailed memories, reading your exact thoughts and knowing your innermost fears and desires? Accessing such data would likely go beyond typical concerns surrounding the notion of privacy.

To explain where I am coming from, let's imagine, for argument's sake, that our minds can extend to some of the technologies we rely on, much like the Extended Mind thesis (Clark and Chalmers 1998) suggests (more on which below). If that's true, it is possible that if others gained unauthorised access to our mind-extending technologies, such as our smartphones or

laptops, they would not be committing a typical violation of privacy—they would be looking into our minds (Reiner and Nagel 2017).

Assuming the above is correct, there are pressing theoretical, ethical, legal and engineering ramifications. In this paper, I want to introduce and outline one of them: Traditionally, in both scientific and legal contexts, it is commonplace to distinguish between two categories of information stored on and generated by our digital devices and cloud accounts: data and metadata. Nevertheless, if the Extended Mind thesis is correct, there is potentially a third kind of digitally stored information that needs introducing—what I will refer to as *mental data*. This new kind of information might require different ethical and legal protection compared to data and metadata and it might call for increased levels of cyber-security.

2. Background

Neuroethics is a fast-growing interdisciplinary approach to applied ethics, whose main focus has been the ethical ramifications emerging from advancements in our ability to understand, read and manipulate *brain activity* (Roskies 2021). For instance, discussions within neuroethics have centred around implications arising from pharmacological enhancements, deep brain stimulation (DBS), electroencephalograms (EEG) and functional Magnetic Resonance Imaging (fMRI), all of which either interfere with or monitor the activity of the *neural* machinery.¹

Despite neuroethicists' focus on the brain, Neil Levy (2007a; 2007b) has suggested that the Extended Mind thesis—which denies that the mind is identical to the brain—“dramatically expands the scope of neuroethics” (2007a, 3). This is an interesting claim with potential ramifications for both the subject matter and the methodology of neuroethics. Nevertheless, Levy has taken advantage of this point only methodologically, by using considerations around the Extended Mind thesis to throw new light on considerations surrounding traditional neuroethical questions, concerning the manipulation and monitoring of the brain (see also (Blitz 2010) for a similar approach). It seems therefore that, despite Levy's remarks, neuroethics has not directed its attention to ethical questions concerning the *non-neuronal* parts of extended minds so far.

¹ See, for example, Levy 2007a, 2007b; Savulescu and Bostrom 2009; Illes and Sahakian 2011; Blitz 2010; Ienca and Haselager 2016; Ienca and Andorno 2017; Lavazza 2018; McCarthy-Jones 2019. For an extensive overview of the field of neuroethics, see Roskies 2021.

Likewise, within the literature on the Extended Mind, there are no sustained ethical treatments covering details of specific, moral questions arising from the possibility of mind extension. Notably, there is a growing body of literature that focuses on the metaphysical question of whether the Extended Mind thesis implies the possibility of extended personhood (e.g., Baker 2009; Wilson and Lenart 2014; Heersmink 2017; 2018; 2020a; 2020b; Harris 2022). This is an important question in the present context, because should the answer be positive, it would indeed have significant ethical ramifications. Additionally, a few authors have suggested that the Extended Mind thesis can potentially have important consequences for several ethical questions, including privacy, autonomy and the choice to cognitively enhance (e.g., Reiner and Nagel 2017; Smart et al. 2017; Heersmink 2017b; Farina and Lavazza 2022). Nevertheless, the details of the far-reaching ethical and possibly legal impact that the Extended Mind thesis can potentially have are rarely discussed.²

In what follows, my aim is to partly rectify this omission by looking into some of the possible ways in which mind extension can alter our perception of important matters concerning privacy. Of course, it is not here possible to singlehandedly give the topic the full treatment it deserves. Hopefully, however, the following will set in motion a much-needed discussion on the way the Extended Mind thesis can impact our understanding of ethical considerations surrounding the concept of privacy, shape the design of relevant technologies, and even inform the corresponding legal framework.

First though, a primer to the Extended Mind thesis is in order.

3. The Extended Mind

In the late 90s, philosophers Andy Clark and David Chalmers introduced the Extended Mind Thesis (Clark and Chalmers 1998)—a controversial but remarkably resilient idea. Essentially, the thesis is a metaphysical claim about the realisation basis of minds: It states that aspects of our minds—including, mental states such as memories, beliefs, desires—extend beyond our organismic boundaries when, to perform certain mental tasks (such as thinking, deciding, planning, or remembering) we *couple* with technological peripherals containing relevant, technologically stored information.³

² For exceptions, see (Carter and Palermos 2016) as well as Milojević's (2017) response to (Carter and Palermos 2016). See also (Drayson and Clark, unpublished).

³ Regarding the metaphysical nature of the thesis, Robert Rupert (an opponent of the view), frames the idea that the mind extends by noting that, at its most basic form, the thesis holds that “human cognitive processing *literally* extends into the environment surrounding the organism, and human cognitive states *literally* comprise—as wholes do their proper parts—elements in that environment; in consequence, while the skin and scalp may

This capacity for *mind extension* is supposed to be the most robust manifestation of our mind's tendency to hook up with technological peripherals (Clark 2003). The more 'primitive' expression of *mind extension* is known as *cognitive extension*, whereby lower-level cognitive processes, such as visual or arithmetical processing, extend to the tools we heavily interact with.⁴ A list of plausible low-tech examples of cognitive extension includes the use of bones as weapons to hunt (depicted nicely by Kubrick as a bifurcation point in hominid evolution in the opening sequence of *2001: A Space Odyssey*), the use of canes by blind people to navigate their environment, and the use of pen and paper to perform complex arithmetical calculations (McClelland et al. 1986).

To focus on the last example from McClelland et al., take a three-digit multiplication problem, such as 987 times 789. Most individuals may only perform such a long multiplication by using pen and paper to externalize the problem in symbols. Once the problem has been transposed from the brain onto the paper using basic arithmetic symbols, the agent can break down the complex problem in a series of elementary tasks, consisting mainly of carrying out simpler multiplications and writing down their results on the paper to use them in later stages. This is a highly co-ordinated process, whereby the eyes and hand of the agent transparently synch to complete the task, with the final, cognitive achievement not being fully reducible to the neural act of reciting the times tables within one's head (Giere and Moffatt 2003). Rather, this is a complex task that involves ongoing, continuous bi-directional interactions between brain, hand, pen and paper, all the while it is being regulated, in the background, by the norms of the underlying notational system—for example, the order in which to perform each multiplication, where on the paper to store the results, how to write them down and so on.⁵

Despite their brow-raising quality, the Extended Mind thesis and the Hypothesis of Cognitive Extension have had, since their inception, a remarkable (perhaps even surprising) impact on several disciplines, ranging over epistemology (e.g., Carter et al. 2014; Carter et al. 2018), computer science (e.g., Smart 2012; 2017), philosophy of science (e.g., Toon 2014), philosophy of emotions (e.g., Krueger and Szanto 2016, Carter et al. 2016), ethics and the law (e.g., Carter and Palermos 2016). Of course, at the same time, the idea that mind and cognition can extend beyond brains has not been without its critics (e.g., Adams and Aizawa 2001, 2008,

encase the human organism, they do not delimit the thinking subject" (Rupert 2004, 389; the emphasis is mine). Likewise, Mark Rowlands (a proponent of the extended mind thesis) writes: "[N]ot all cognitive processes occur inside the head or skin of cognizing organisms. Therefore, if we assume that the mind of a cognizing organism such as a human being is made up, at least in part, of cognitive processes, the central *metaphysical assertion* of this book is that the mind is not, exclusively, inside the head" (Rowlands, 1999, 29; the emphasis is mine).

⁴ Clark and Chalmers (1998) are the first to distinguish between the Extended Mind Thesis and the Hypothesis of Extended Cognition.

⁵ For details on the effects of external representational systems and of their normative features on cognition, see Menary (2007).

2010; Rupert 2004, 2009): Claiming, after all, that mind and cognition are not necessarily brain-bound challenges many people's intuitions and goes against centuries of brain-centrism (i.e., the default position that our minds just are our brains).

My intention here is not to offer an extensive defence of the Extended Mind thesis. Nevertheless, it is important to work with a plausible take on it. To this end, I will quickly address what is perhaps the most obvious (and pressing) worry against the view: Assuming that the thesis is true, such that our minds do sometimes extend, how do we know when a tool constitutes part of our mind and when not? After all, there are multiple technological devices we employ on a regular basis: Does my dishwasher count? What about my desk lamp?

A rough and ready response offered by people working on the topic (e.g., Clark and Chalmers 1998, Chemero 2009, Krueger 2012, Palermos 2014) is to focus on the concept of *coupling*. Two systems—for example, a human and an artifact—are coupled and thereby constitute a single integrated system, when, in the process of performing some task, they bidirectionally affect each other on a continuous basis.⁶ Following this line of thought, an agent's mind or cognition are extended, when, in performing a mental or cognitive task, the agent is coupled with a tool by bidirectionally interacting with it. Dishwashers and desk lamps are not therefore devices that can support the extension of mind or cognition, because, when we employ them, the causal effect is one way—from the tool to the agent but not *vice versa*. In other words, while we rely on tools such as desks and desk lamps to perform mental and cognitive tasks, when we do, we do not continuously and bidirectionally interact with them. Therefore, there is no coupling and so (in line with intuition) minds or cognition do not extend to them.

A blind person's cane while they are navigating their environment, however, or (as the above description indicates) pen and paper while performing long multiplications are better candidates in this respect. More plausible, perhaps, is the employment of smart phones while we use their SatNav functionalities to find our way around a new city. Better still is the case of Tactile Visual Substitution Systems (see, for example, Bach-y-Rita and Kercel 2003, Kiverstein and Farina 2012, Kiverstein, Farina and Clark 2015).

Roughly, a Tactile Visual Substitution System (TVSS) is made up of a mini video camera (sometimes attached to a pair of glasses) and a unit that transforms the visual input from the camera into tactile stimulation under the agent's tongue, or on the skin of their forehead (or their back). By moving around, agents affect the input to the camera, which affects the tactile

⁶ The notion of a coupled system comes from the branch of theoretical mathematics known as Dynamical Systems Theory.

sensations they receive, which, in turn, determines their subsequent motions, and so on. On the basis of this densely interactive process between organism and machine, blind agents quickly start perceiving shapes and objects and orienting themselves in space. Occasionally, they also describe their experience in quasi-visual terms (Macpherson 2018), indicating that they are enjoying phenomenal qualities close to those of the original sense modality being substituted. In such cases, neither the biological organism nor the technological artifact can see; but the extended integrated system that is established on the basis of their continuous mutual interactions is capable of quasi-visually perceiving the world.

According to the above, then, the use of canes, pen and paper, SatNavs and TVSSs are plausible candidates of cognitive extension, whereby an agent's cognitive processes of navigating in space, performing arithmetical calculations or quasi-visually perceiving the world extend beyond the agent's organismic boundaries. But, shifting our attention away from *cognitive* and toward *mind* extension, we may also ask: Do existing apps—such as the photo-taking-and-categorising apps, e-calendars, and note-taking apps of our smartphones—with their increasing capacity to contribute to what seems to be the encoding, storage and retrieval of detailed memories, plans, beliefs and desires, also count as *mind extensions*?⁷ The answer, following the above line of thought, seems to largely depend on how we use these apps: Can we and do we couple (by bidirectionally interacting) with them, while performing *mental* tasks throughout our daily activities?

This question—concerning our minds' coupling with recent and current technologies—is an important one. Admittedly, peoples' intuitions, when asked to comment on it, often pull toward different directions. On my part, I suspect that each persons' stance toward the Extended Mind thesis largely depends on their own experience and cognitive strategies—do they often perform mental tasks by bidirectionally interacting with existing technologies? If the answer is positive, they are more likely to feel the force of the thesis. Nevertheless, let's not lose sight of the larger picture here: Whatever the status of *existing* technologies may be, one can easily imagine that, in a few years' time, when brain-machine interfaces (e.g., Elon Musk's Neuralink) will have been perfected, the coupling between our brains and technologically stored information will be as seamless as the one we currently entertain with the data bathing in our wetware. When that is the case (and it may not be too long before it is)—i.e., when we command and retrieve information from our smart devices merely by thinking about it—it will be increasingly hard for anyone to deny that such devices and the information they contain are

⁷ See also Heersmink (2017a) for interesting discussion of the way lifelog technologies might extend the capacity of autobiographical memory.

proper parts of our minds. In fact, even the view's fiercest critics are quick to note that, while they doubt its truth so far, it is possible that the Extended Mind thesis may materialise in the future.⁸

Now, to take stock, it was noted that the goal here is not to argue for the Extended Mind thesis or any specific instances of it. Instead, the claim is that *if* the Extended Mind thesis is true—if, indeed, the mind can extend beyond our organismic boundaries—then this is most plausibly the case, when (and only when), to perform some mental task (such as thinking, remembering, planning, deciding) we *bidirectionally interact*—i.e., *couple*—with certain aspects of the environment. Put another way, the presence of coupling may act as the rule of thumb by which we can judge whether an external portion of the environment and the information it contains may count as part of our minds (assuming, of course, that the Extended Mind thesis is true).⁹

But then here is a worry: If we accept that this plausible take on the Extended Mind thesis is true, such that minds do extend to coupled environmental aspects containing information that may count as our memories, beliefs, desires and, more generally, as *mental content*, then, as Reiner and Nagel (2017) also note, it transpires we are exposed to a threat that has so far been thought impossible: Suddenly, the detailed contents of our minds are (or at least will soon be) literally open for others to scrutinize.

⁸ Adam and Aizawa (2001) describe themselves as ‘contingent’—rather than necessary—‘intracranialists’ about cognition, and Rupert (2004, 394) notes: “I take for granted the weaker modal claim, that extended cognition is possible.”

⁹ A further clarification might be helpful here: Notably, to make their argument for the Extended Mind thesis, Clark and Chalmers (1998) claim that when agents have reliable, easy access to peripherals (such as diaries and life-logging technologies), whose contents can be automatically endorsed, these contents partly constitute dispositional (i.e., non-occurrent) beliefs. Thus, Clark and Chalmers claim, dispositional beliefs, and, thereby, minds extend. I am not entirely convinced of the success of this line of argument for the Extended Mind thesis (for details as to why, see also (Palermos 2014; 2018)). As the above indicates, however, it is possible to argue for the claim that memory and mind extend in an alternative way. According to this alternative approach, mind extends not because dispositional beliefs extend, but because the whole mental process of memory—i.e., encoding, storing and retrieving information—extends, when (and only when) agents are coupled to technological peripherals. In such cases, we may speak of an extended mind, to which the corresponding extended mental process belongs.

There is of course the further question of whether one should call the externally stored informational contents of this extended memory system ‘memories’ or ‘dispositional beliefs’ (for reasons to doubt that we should, see Weiskopf 2010). But, if we accept that the mind extends in such cases (because the mental capacity to remember extends), little doubt should there be that any externally stored informational contents that represent what we usually refer to as memories, thoughts and desires should count as *mental contents*. Personally, I am content (and I will continue) to loosely refer to such externally stored mental contents as ‘memories,’ ‘thoughts’ and ‘beliefs’—if anything, and as the debate over eliminativism suggests (e.g., Churchland 1981, Dennett 1987, Egan 1995), no one really knows what these terms really refer to, and we may come to discover that they refer to nothing specific at all. But whatever the case may turn out to be, my preferred, alternative line of argument for the Extended Mind thesis does not rest on the validity of using such folk-psychological terms for referring to the external portions of the mind.

How should we respond to this existential threat (to use a repurposed, yet popular in the context, philosophical term)?

4. The Data | Metadata Distinction

Discussions on the right to privacy in the digital age have mainly focused on the conditions under which it is legal to generate, retain or process *metadata* in the absence of users' consent. Standardly, the notion of metadata is understood by contrast to the notion of data. The distinction between the two has existed for several decades. It was initially drawn to guide electronic surveillance law pertaining to telecommunications; at first, communications via the telephone and later via email. Roughly, data (or "content") referred (and still refers) to the actual communicated information during a communication event (what the interlocutors discussed or wrote to each other), whereas metadata (or "non-content") is used to refer to information *about* the communication event: for example, that a certain type of communication took place, the time it took place, the media that were involved, how long the communication lasted and so on (for more details, see (Bellovin et al. 2016).

Thus, with that distinction in place, the situation, so far, has roughly been as follows: Communication data has been illegal to obtain without the individual's consent, unless the enforcement of law demands it. In contrast, the collection, storage and use of metadata may sometimes lawfully take place in the absence of the related person's consent, a contractual agreement with them, the need to comply with some legal obligation or the demands of law enforcement.

This may initially sound unproblematic—and perhaps it had been for several decades. Nevertheless, several authors have recently commented that, from a legal perspective, current and emerging technological advances have rendered the distinction between data and metadata, as well as the associated laws, outdated. Steven Bellovin, Matt Blaze, Susan Landau and Stephanie Pell (2016) note, for example, that

1. The concept of metadata, as a category of communication information that is wholly distinguishable from communications content, is outdated;
2. The traditional physical and legal distinction between content and non-content, which has generally provided a consistent, reliable method for discerning more sensitive or revealing aspects of communication information worthy of Fourth Amendment protections is too difficult to apply;
3. The application of traditional content/non-content distinctions leads to inconsistent and anomalous results

As the authors note further, a particularly important reason for this unfortunate situation is that some information may fit into either category: “Because the content of a communication can sometimes be inferred from its corresponding metadata [...] it is not clear that distinct, meaningful legal lines can be drawn between these two categories of information in the way it could be done during the telephony era” (ibid., 92). Thus, the authors conclude, “the concept of metadata as a category of information that is entirely distinguishable from communications content and thus deserving of lower privacy protection is no longer tenable” (ibid., 92).

This is, no doubt, an important topic, requiring urgent attention. Nevertheless, I do not intend to here delve into the details of the data/metadata distinction any deeper. Rather, for ethical and legal purposes, I want to suggest, an additional category of data might need to be introduced—*viz.*, ‘mental data.’

5. Mental Data

It was argued, in section 3, that if the Extended Mind thesis is true, then this is most plausibly the case when (and only when) in the process of carrying out mental tasks, agents bidirectionally interact with artifacts, effectively coupling (i.e., integrating) with them. Think of cases where people mutually interact in such dense a way with photo-taking-and-categorising apps, e-calendars, note-taking apps, reminders apps, in order to perform such mental tasks as storing and recalling autobiographical information, planning their schedules and organising their life, reaching complex decisions about their work, storing and recalling semantic information that is closely relevant to themselves and the roles they identify with.¹⁰ If the Extended Mind thesis is correct, then such cognitively integrated applications count as proper parts of their minds, and the information they contain may count as *mental* content. In other words, any data stored on the cognitively integrated applications of technological peripherals may count as *mental data* (assuming, of course, that the Extended Mind thesis is true).

Now, if the above is correct—I want to submit for further consideration—mental data is a distinct kind of data category that requires unique treatment from both an ethical and a legal perspective. Such a suggestion, of course, raises several questions: How is mental data different from and how does it relate to data and metadata? How might mental data be

¹⁰ For interesting discussion on the way ‘life-logging’ technologies may extend one’s autobiographical memory, see (Heersmink 2020a; 2017a).

protected in ways that are different from the ways data and meta-data are protected, and what considerations might motivate such difference in treatment?

I turn to these questions right away.

6. Differences, Similarities, Overlaps

Let's start with the distinction between data and mental data, which is relatively straightforward. Within the legal context, the term 'data' or 'contents' has been traditionally used to refer to the contents of electronic communications: for example, what was discussed over the phone, what the body of an email said, what one texted their friend. Thus, one way that mental data is different to data is that the former refers to information that is stored on devices yet without forming part of any electronic communication. Moreover, to count as mental data, the relevant informational content must form part of an application or device with which the agent bidirectionally interacts in the process of completing mental activities.

Turning now to the notion of 'metadata,' the term has been used in the context of law to refer (as the term indicates) to data *about* data. This may be taken to imply, given what's mentioned above, that 'metadata' concerns only electronic communication data. However, as a moment's reflection may suggest, any electronic transaction can potentially (and often does) generate metadata—from offering positive audience feedback (i.e., liking posts) on social media to making electronic financial transactions online or at an ATM. Arguably, and despite their ubiquity, metadata are always distinct from mental data. For example, taking a photo and storing it on your smartphone as means for encoding, retaining and retrieving (i.e., memorising) autobiographical information may generate metadata such as the fact that you took a photo, your location when the photo was taken, how big the file containing the photo was, etc. But the actual informational content of the photo is expressly distinct from this other information *about* the photo. Likewise, using a notes-taking app to write some of your thoughts down while working on a project might generate metadata regarding the fact that you used the app, what time you used the app, etc. But your note's content is obviously different from this information *about* the note.

Nevertheless, the distinction between mental data and metadata might not always be as clear-cut. As is often highlighted, accessing one's metadata in bulk may reveal significant portions of one's mind. For example, the 2014 United Nations General Assembly resolution notes that "certain types of metadata, when aggregated, can reveal personal information and can give an insight into an individual's behaviour, social relationships, private preferences and

identity” (UN General Assembly 2014).¹¹ Indeed, today there are specialised AI algorithms dedicated to performing psychometric analyses of individuals based on the metadata generated by their online transactions. As Uwe Peters (2022) notes:

AI systems for psychometric tracking are now employed by, for instance, social media companies and consultant firms (see the Cambridge Analytica scandal) to segment website users into personality groups for tailored advertising and messaging (Hinds and Joinson 2019). Indeed, IBM offers a program (‘IBM Watson Personality Insights’) to private individuals for predicting website users’ ‘personality, needs, and values’ based on social media data for highly targeted messaging.

Therefore, there seem to be good reasons for suspecting that the distinction between metadata and mental data is not straightforward and that, most certainly, it is a difficult one to draw. Of course, one might attempt to retain the distinction between the two by pointing out that one’s character and personality traits are no more than mere *attributes*—very much unlike proper parts—of one’s minds.¹² Accessing, therefore, metadata capable of revealing this kind of information may not be quite the same as accessing one’s mental data. Or, perhaps, an alternative way to put the same point would be to note that there is a principled distinction to be drawn here based on notions such as ‘constitution’ and ‘representation.’ Metadata (an argument like this would go) can only *represent* rather than *constitute* one’s mind, such that accessing one’s metadata in bulk is not the same as accessing one’s mental data. Such objections, it should be clear, however, would themselves require considerable defence.

It is therefore unclear whether the notions of metadata and mental data may—at least sometimes—overlap. The preceding discussion has done little more than merely scratching the surface of the issue, but it is important to note here that this is potentially a pressing topic, worthy of serious consideration. The reason, as I am about to argue, is that data, metadata, and mental data may require significantly different levels of legal and even technological protection.

7. Levels of Protection

It has been argued, so far, that if the Extended Mind thesis is true then some of the informational content that is stored on devices and apps with which we bidirectionally

¹¹ In his role as *Amicus Curiae* in support of the lawsuit that called into question the lawfulness of the National Security Agency’s bulk collection of phone metadata, Michael Lynch also attempts to condemn the collection of metadata by noting that it could harm the privacy of our minds.

¹² A good reason one might wish to deny that metadata and mental data overlap has to do with personalisation of services. If certain metadata were to qualify as mental data—thereby requiring, as I am about to argue, greater levels of protection—it would impose significant restrictions on how websites and applications may generate, retain, and make use of them, potentially resulting in less personalised services.

interact (so as to complete mental tasks) count as mental data. Such informational contents, in other words, are proper parts of our minds. If correct, this claim raises important privacy concerns.

For the most part, our biological minds are private to a near absolute degree. Of course, this is not to deny, that, in certain situations, others may be able to partly infer some of what we are thinking or feeling.¹³ Nevertheless, the *detailed contents* of our minds are private and inscrutable. Others are unable to read, hear, see or in any other way come to know another's memories, thoughts and desires in the full detail that has so far been *exclusively* available to the individual to whom these memories, thoughts and desires belong. And, crucially, this is so no matter what the law may demand. If I am sufficiently determined, my detailed thoughts, memories and ideas are mine and mine alone—unless, of course, I decide to voice them. And if I also put my poker face on, the same is true of my emotions and desires too. No interrogation technique or law can make me reveal the detailed contents of my mind if I choose not to.

But then here is the issue: If the detailed contents of our biological minds have so far been impossible for others to obtain in the absence of our will, and if we do value mental privacy as we have so far entertained it, then the same should remain true of the contents of our extended minds. Mental data, that is, should entertain the same degree of privacy as the detailed contents of our brains.

Why require the *same* degree of privacy between these two physically distinct kinds of data? The answer, again according to the Extended Mind thesis, is that mental data, though physically realised by non-biological processes, they are ontologically the same as the contents of our brains: That is, metaphysically speaking, both kinds of data are part of the mind and thereby constitute mental contents—with their actual material basis and spatial location making no difference to their ontological status. In other words, provided that the Extended Mind thesis is true, mental data are essentially no different to the data stored in the neural web of our brains; consequently, access to mental data (i.e., data stored on devices that have been cognitively integrated) would be the same as having a third entity peeking into and decoding one's brain's activity. Thus, to preserve mental privacy as we know it, mental data should be equally inaccessible to the detailed contents of our brains.

What is more, the value of mental privacy is not the only value that could motivate this normative claim. The values of *mental autonomy* and *psychological individuality* can play

¹³ If the 'direct perception' account of social cognition is correct ((Gallagher 2008; Gallagher and Hutto 2008)), we may even have "direct perceptual access to aspects of another's mentality" (Krueger 2012, 150).

important roles in its support as well. With regards to mental autonomy, for example, Simon McCarthy-Jones (2019) notes that “exposure of our thoughts would effectively alter them by pressurizing us not to think certain things.” Specifically, “if we cannot keep our thoughts private then we will experience conformity pressures to think in a certain way.[...] This pressure will cause self-censorship, impairing mental autonomy.” This interesting idea—that mental censorship would cripple mental autonomy—seems to be also echoed in a related point made by Blitz (2010). Blitz notes that even though international and constitutional law has placed significant emphasis on the right to freedom of speech, opinions expressed in past rulings of the U.S. Supreme Court suggest that this well-known right might, in fact, be part of a broader and more fundamental concept—that of ‘*freedom of mind*’ or ‘*freedom of thought*.’ One plausible explanation for this claim, we can here note, is to point out that the things of which one speaks shape the thoughts one may entertain, which, in turn—through public discourse—can be criticised, re-evaluated and developed further.¹⁴ Hence, from this point of view, the right to freedom of speech may rightly be seen (as Blitz (2010, 2) notes, referencing *Wooley v. Maynard*, 430 U.S. 705) as a component of the more important right to freedom of mind or thought (or, of what McCarthy-Jones (2019) refers to as the right to mental autonomy, which, itself, is supposed to be an important aspect of freedom of thought). But then, following the same line of reasoning, it can be likewise argued that the right to mental privacy is also an essential part of our fundamental right to freedom of thought:¹⁵ Quite obviously, censoring one’s thoughts and mental contents is at least as damaging to one’s freedom of thought as censoring one’s words; and, quite likely, even more detrimental, given the direct nature to the *mental* censorship involved. Thus, in line with McCarthy-Jones’ point, it seems reasonable to hold that if mental autonomy and freedom of thought are valued, mental privacy should be thoroughly protected.

Or consider individuality: Defending the unlawfulness of metadata collection in bulk, Michael Lynch (2013) notes:

The connection between informational privacy and autonomous personhood have deep roots in the very nature of human psychology and therefore in our very conception of ourselves as personal beings. As noted above, part of what makes your individual mind *your* mind is that you have a degree of privileged access to your mental states—the various thoughts and feelings that comprise your individual mind. And that includes, crucially, the ability to control

¹⁴ The idea that language (both written and spoken) can actively drive thinking can be easily framed in terms of the extended mind thesis. See, for example, Clark (1998).

¹⁵ Ienca and Andorno (2017) also suggest that mental privacy is an important component of freedom of thought.

access to the content of those thoughts and feelings; that is, to choose whether and when one shares this information with others.¹⁶

Lynch's points, which he raises to promote the protection of metadata, should be all the more plausible when we consider them in relation to protecting mental data. Indeed, little doubt should there be that an important aspect of what it means to be a psychological individual is to possess the ability to stand as a separate psychological entity—an ability that seems heavily predicated on possessing mental borders, which (especially absent one's consent) are for the most part impenetrable.¹⁷ Thus, if others started having access (and especially in a non-consensual manner) to the detailed contents of our minds (e.g., our beliefs, memories, thoughts and desires)—informational contents to which, so far, we alone had direct access—then, in an important sense, we would no longer stand as the separate, individual, psychological entities that we currently take ourselves to be.

Valuing therefore mental privacy in itself or for its contribution to mental autonomy and psychological individuality may act as a strong drive for providing mental data with increased levels of protection. In fact, attempting to raise the protection of data that may qualify (in the present terminology) as mental data is not entirely without precedent in the legal context. Think, for example, *Riley vs. California* (573 U.S. 373 (2014)), which revolved around the question of whether the Search Incident to Arrest (SITA) applies to smartphones. According to the court ruling's syllabus, SITA allows police to carry out—without a warrant—a search of the area that falls within the immediate control of a person that is under arrest (including, for example, their pockets, bag or car). In the written opinion of the United States Supreme Court, Chief Justice John Roberts notes first that “modern cell phones, [...] are now such a pervasive and insistent part of daily life that the proverbial visitor from Mars might conclude they were an important feature of human anatomy” (p. 9). So, the report concludes...

the fact that technology now allows an individual to carry such information in his hand does not make the information any less worthy of the protection for which the Founders fought.

¹⁶ In support of his claims, Lynch cites Edward Bloustein: “The fundamental fact is that our Western culture defines individuality as including the right to be free from certain types of intrusions. This measure of personal isolation and personal control over the doctrines of its abandonment is of the very essence of personal freedom and dignity, is part of what our culture means by these concepts” (Bloustein 1964, 973).

¹⁷ As Lynch (2013) also notes invoking Descartes' *Meditations on First Philosophy* and Locke's *Essay Concerning Human Knowledge*, “a person's privileged access to his or her thoughts is not only the most basic and original form of informational privacy, it is part of what differentiates one person from another and thus helps to constitute our status as individuals.”

Our answer to the question of what police must do before searching a cell phone seized incident to an arrest is accordingly simple—get a warrant (p. 28).¹⁸

It is reassuring to see that courts are willing to raise the level of protection entertained by technologies with which we cognitively integrate. The crucial question, however, is whether existing law offers sufficient levels of protection. In the case of *Riley vs. California*, for example, the court seems to think that searching into the contents of one's smartphone should be acceptable if a warrant has been obtained. This raises the important question of whether protection of our mental data should be *qualified* or *absolute*: Should there be exceptions to whether mental data can be lawfully accessed by a third party, such that warrants to do so can be issued? As McCarthy-Jones (2019) notes,

It is possible that society may wish permissible violations to exist. Warrants can be sought and granted to search individuals' homes. It is not immediately obvious that there should not be similar permissible violations for the mind. It is possible to imagine a judge granting a warrant for a mental search.

Most authors (e.g., Ienca and Andordno, 2017; McCarthy 2019) who have already contemplated on the level of privacy that our minds should entertain have only considered levels of *legal* protection;¹⁹ i.e., whether protection of mental privacy should be qualified or absolute. Here, I want to highlight the additional possibility of *technological protection* with regards to mental data. That is, in addition to the protection that the law may offer, further technological options might in fact be available—options which may offer even stronger levels of protection than absolute legal protection.

Why might we wish to offer mental data stronger protection than what absolute legal protection may offer to them? The answer is that, so far, the detailed contents of our minds have been practically, not just legally, inaccessible to others. Moreover, as the previous discussion indicates, this practical inaccessibility of the detailed contents of our minds seems to have also distinctively shaped our sense of individuality and contributed to our mental autonomy. Therefore, if mental privacy has so far meant that detailed mental contents are practically inaccessible, and if we do value mental privacy as we know it—either inherently or instrumentally, because of the way it has contributed to mental autonomy and individuality—

¹⁸ The court's written opinion is available here: https://www.supremecourt.gov/opinions/13pdf/13-132_8!9c.pdf

¹⁹ These authors focus primarily on the separate issue concerning access to brain data by means of fMRI and EEG techniques. The only possible exception here is Lavazza (2018), who also focuses on the protection of brain data (rather than mental data), but with the use of technological means.

then mental data ought to remain practically—not just legally—inaccessible. Put another way, when it comes to mental data the question is *not whether or when* they should be legally accessible but whether they should be accessible *at all*. This is a meaningful question to ask, because existing encryption technologies can ensure that third party access to sensitive data can be made practically impossible.²⁰

What this means, practically and theoretically, is that there is a menu of protective options that data can be offered, depending on which category of data they fall under. Listing them from weakest to strongest kinds of protection, we may choose, for example, from the following (non-exhaustive) list:

1. Qualified legal protection
2. Absolute legal protection
3. Technologically impossible to obtain, absent the agent's consent
4. Technologically impossible to obtain

Data and metadata, as we have seen, have so far only entertained varying strengths of relative legal protection. Ienca and Andorno (2017) have suggested that brain data (obtained via fMRIs or EEGs) should, perhaps, be legally protected in an absolute manner. McCarthy-Jones (2019) has suggested the same about internet searches, notebooks and diaries,²¹ and Lavazza (2018) appears to suggest option 3 with regards to brain data. So far, however, no one has considered options 3 and 4 with regards to data that may qualify as mental data.

And yet, given the preceding analysis, this is perhaps precisely the way to go. Specifically, the idea put forward for consideration is that data which, on the Extended Mind thesis, would qualify as mental data, should be either technologically impossible to obtain or technologically impossible to obtain, absent the agent's consent.²² In other words, although data and metadata have been so far protected by legal means, options 3 and 4—which appeal to existing technological resources—should now be considered in relation to mental data, as

²⁰*Practically* impossible, because, given that there are no cryptographic techniques that are immune to cryptanalysis techniques—save the One Time Pad (OTP), which works only under very strict, practically impossible to satisfy conditions—any data are in theory accessible. Existing encryption methods, however, can ensure there are no efficient (i.e., non-extreme-time-consuming, brute force) methods to access encrypted data.

²¹ Interestingly, McCarthy-Jones suggests this by appealing to the Extended Mind thesis, though, unfortunately, he is not going into sufficient detail. Although I am open to accepting the claim (depending on the specifics of how it is understood), I am not entirely convinced that McCarthy-Jones' argument for the claim succeeds.

²² Likewise, metadata that are likely to qualify as mental data (see the discussion of section 6) should, perhaps, also be technologically impossible to obtain, absent the agent's consent. Such a measure would sufficiently protect user's mental privacy, while also providing them with the ability to opt-in for personalised services (even if in the expense of some of their mental privacy). See also fn. 10.

the means to match the level of privacy that our skulls have so far offered to our detailed thoughts, memories, feelings and desires.

Nevertheless, an additional, cautionary note regarding the choice between 3 and 4 is in order here. Legally adept readers may note that when consent is present, the corresponding data would be subject to the ‘third party doctrine,’ according to which individuals who volunteer information to another party do not have, in relation to the voluntarily shared information, any “reasonable expectation of privacy.”²³ This likely indicates that the third party doctrine might itself need to be revised and refined in view of the fact that consent might now be provided for information that is so sensitive as to qualify as mental data.²⁴ Additionally, when, instead of option 4, option 3 is employed, a further question to consider concerns whether it should be incriminating or even legally non-neutral to deny access to one’s mental data, given one is in a position to grant it. Here, considerations regarding the value of dignity and the right to protect one’s dignity might suggest that refusing access to one’s mind (or parts of it) should come with no repercussions—though, to avoid placing users in such predicament altogether, technologies should, perhaps, be designed so that users can always choose to opt for option 4, with regards to at least some of their mental data.

8. Discussion and Conclusion

The overall argument offered for consideration proceeds in the following way:

- P1 - *Mental Privacy*: So far, the detailed contents of our minds have been *practically*, not just legally, inaccessible to others.
- P2 - *Extended Mind*: If the Extended Mind thesis is true, then the informational contents of the apps and devices with which we bidirectionally interact in the process of completing mental tasks qualify as mental contents, i.e., *mental data*.
- P3 - *Value*: If we value mental privacy as we know it—either inherently, or instrumentally for its contribution to psychological individuality and mental autonomy—then we should protect it.

²³ For details on the notion of “reasonable expectation of privacy”, see the concurring opinion of Justice John Marshall Harlan on *Katz vs. United States*, 389 U.S. 347 (1967).

²⁴ For more details on how the ‘third party’ doctrine creates difficulties and raises questions for its ability to protect (in its current form) sensitive metadata that have been shared via internet technologies, see (Bellovin et al. 2016). Wittes and Chong (2014) also offer interesting discussion on the issues raised by the third-party doctrine with regards to the data and metadata that are bound to be generated in a society where the use of mind-extending technologies is ubiquitous.

C: If the Extended Mind thesis is true and if we value mental privacy as we know it, mental data should be not just legally protected, but practically inaccessible to others (unless, perhaps, the agent consents in sharing them).

Simply put: If the Extended Mind thesis is true and if we value mental privacy as we know it, then the informational contents of our extended minds should be practically inaccessible by others—unless, perhaps, the agent is willing to share mental data themselves.

While this was not the place to argue for the Extended Mind thesis, I have partly defended premise 2 by offering a plausible take on the thesis with the aim of dealing with what has come to be known as the ‘cognitive bloat’ worry (Rowlands 2009) (i.e., the objection that if we do not have a principled way to delineate what may count as proper part of one’s extended mind, then too many things will qualify, leading to a *reductio ad absurdum*). Briefly, my claim has been that a device or application and its informational contents count as part of one’s mind if and only if the agent bidirectionally interacts with it to complete mental tasks.

With regards to P1, though not indubitable, I have assumed that it is largely correct. While it is true that, sometimes, others may be in a position to roughly infer some of what we are thinking, feeling, desiring, and may even come to directly perceive aspects of our mindedness (as per the direct perception account of social cognition), it is largely impossible for them to come to know our thoughts, memories, feelings and desires in the full detail that is exclusively available to us.

By contrast, P3, which concerns a value question, might seem more questionable, and people may respond to it in different ways. Some, perhaps, would be willing to give up mental privacy (at least as we know it), alongside the associated values of individuality and mental autonomy. Several reasons might motivate this stance. One of them concerns user safety. Future technologies may require that experts gain access to users’ mental data to repair damages with potentially adverse psychological effects. This issue could be overcome, however, by providing users with the ability to opt for the third kind of protection (practical inaccessibility, absent the agent’s consent) rather than the first (mere qualified legal protection—see also section 7).

Another issue that could speak against retaining mental privacy as we know it concerns security. State agencies and intelligence services could make significant advances against terrorism if they could gain access to the mental data of suspects for terrorism. Doing so, however, would require protecting mental data only legally and only in a

qualified manner. Therefore, it makes sense to expect that many people would be hesitant to condone this attitude and that privacy activists would vehemently oppose it.

Finally, curiosity or wander could be a further drive for abandoning mental privacy. It is not inconceivable that a few would like to experiment by transcending the barriers of the mind and individuality, eager to experience and explore ‘telepathic,’ direct, mind-to-mind communication. Thankfully, however, this admittedly enticing drive for giving up mental privacy might only require option 3—i.e., technological protection whereby accessibility to one’s mental data will be impossible, absent the agent’s consent.²⁵

We see, therefore, that various reasons exist for justifying access to others’ mental data, and this to various degrees. Moreover, simply to insist that mental privacy should remain practically absolute, because so far it has been this way, would come suspiciously close to committing the naturalistic fallacy. As I have pointed out, however, the choice to retain mental privacy, as we know it, does not have to be the result of blind adherence to what nature has dictated so far. Rather, mental privacy seems to also be intricately connected with the rather weighty values of mental autonomy and individuality. Thus, given these stakes, it should be obvious that, at the very least, caution and a healthy dose of conservatism should be exercised. Radically altering the nature of our minds and psyches, without first contemplating on the possible individual and social consequences would be, to put it mildly, irresponsible.

Of course, to finally come to an end, I do not presume to have here offered an exhaustive analysis of the issue at hand. Hopefully, however, the notion of mental data, as it has been here introduced, will draw the attention of experts from various disciplines, including ethics, philosophy of law, computer science, philosophy of mind, philosophy of technology. Interdisciplinary research can help us critically evaluate the merits of the present proposal and, in the process, provide us with the chance to carefully consider whether our next technological steps will be in the right direction.

Acknowledgements

I am thankful to two anonymous referees for *AJOB Neuroscience* for their insightful comments, which helped me prepare the final version of the paper. I am also helpful to Alex Dietz, Panos Paris, and the audiences of the visiting research seminar at the University of Patras, the Web

²⁵ It worth keeping in mind, however, that one obvious problem with the option of technological protection, absent consent, is that we often give consent all too easily—do we ever read to what we consent online?

and Philosophy Workshop (part of WebSci'21) and The Society for Philosophy and Technology Conference 2021 for their valuable feedback on previous versions of this paper.

References

- Adams, F., & Aizawa, K. (2008). *The Bounds of Cognition*. Blackwell Publishing Ltd.
- Adams, F. and Aizawa, K. (2010). 'Defending the Bounds of Cognition'. In *The Extended Mind*. (2010), Menary (ed.) Cambridge, Massachusetts, MIT press.
- Adams, F., & Aizawa, K (2001). 'The bounds of cognition'. *Philosophical Psychology*. Vol. 14, No.1, 43-64.
- Bach-y-Rita, P., & Kercel, S. W. (2003). Sensory substitution and the human-machine interface. *Trends in Cognitive Science*, 7(12), 541–546.
- Bloustein, E. (1964). Privacy as an Aspect of Human Dignity, *N.Y.U. L. REV.* 39, 962.
- Carter, J. A., Clark, A., Kallestrup, J., Palermos, S. O., & Pritchard, D. (Eds.). (2018). *Extended epistemology*. Oxford University Press.
- Carter, J. A., Kallestrup, J., Palermos, S. O., & Pritchard, D. (Eds.). (2014). Extended Knowledge. Special issue of *Philosophical Issues*, Vol. 24 (1): 1-482
- Carter, J. A., Gordon, E. C., & Palermos, S. O. (2016). Extended emotion. *Philosophical Psychology*, 29(2), 198-217.
- Carter, J. A., & Palermos, S. O. (2016). Is having your computer compromised a personal assault? The ethics of extended cognition. *Journal of the American Philosophical Association*, 2(4), 542-560.
- Baker, L. R. (2009). Persons and the extended-mind thesis. *Zygon*, 44(3), 642-658.
- Bellovin, S. M., Blaze, M., Landau, S., & Pell, S. K. (2016). It's too complicated: How the internet upends Katz, Smith, and electronic surveillance law. *Harv. JL & Tech.*, 30, 1.
- Blitz, M. J. (2010). Freedom of thought for the extended mind: Cognitive enhancement and the constitution. *Wis. L. Rev.*, 1049.
- Chemero, A. (2009). *Radical Embodied Cognitive Science*. MIT press.
- Churchland, P. M. (1981), "Eliminative Materialism and the Propositional Attitudes", *Journal of Philosophy* 78: 67-90.
- Clark, A. & D. Chalmers (1998). 'The Extended Mind', *Analysis*, 58: 10–23.
- Clark, A. (2003)., *Natural-Born Cyborgs: Minds, Technologies and the Future of Human Intelligence*. Oxford University Press.

- Clark, A. (1998). Magic words, how language augments human computation. In P. Carruthers & J. Boucher (Eds.), *Language and thought: Interdisciplinary themes*. Cambridge: Cambridge University Press
- Dennett, D. C. (1987), *The Intentional Stance*. Cambridge, MA: MIT Press.
- Descartes, R. (1986). *Meditations On First Philosophy* (John Cottingham trans., Cambridge University Press)
- Drayson, Z. and Clark, A. (unpublished). Augmentation, agency, and the spreading of the mental state.
- Egan, F. (1995). Folk psychology and cognitive architecture. *Philosophy of Science*, 62(2), 179-196.
- Farina, M., & Lavazza, A. (2022). Incorporation, Transparency and Cognitive Extension: Why the Distinction Between Embedded and Extended Might Be More Important to Ethics Than to Metaphysics. *Philosophy & Technology*, 35(1), 1-21.
- Gallagher, S. (2008). Direct perception in the intersubjective context. *Consciousness and Cognition*, 17(2), 535–543.
- Gallagher, S., & Hutto, D. D. (2008). Understanding others through primary interaction and narrative practice. In J. Zlatev, T. P. R. Racine, C. Sinha, & E. Itkonen (Eds.), *The shared mind: perspectives on intersubjectivity* (pp. 17–38). Amsterdam: John Benjamins Publishing Company.
- Giere, R. N., & Moffatt, B. (2003). Distributed cognition: Where the cognitive and the social merge. *Social studies of science*, 33(2), 301-310.
<https://doi.org/10.1177/030631270303332017>
- Harris, K. R. (2022). Why the Self Does Not Extend. *Erkenntnis*, 87, 2645–2659.
- Heersmink, R. (2020a). Extended mind and artifactual autobiographical memory. *Mind & Language*.
- Heersmink, R. (2020b). Varieties of the extended self. *Consciousness and Cognition*, 85, 103001.
- Heersmink, R. (2018). The narrative self, distributed memory, and evocative objects. *Philosophical Studies*, 175(8), 1829-1849.
- Heersmink, R. (2017a). Distributed selves: Personal identity and extended memory systems. *Synthese*, 194(8), 3135-3151.
- Heersmink, R. (2017b). Extended mind and cognitive enhancement: Moral aspects of cognitive artifacts. *Phenomenology and the Cognitive Sciences*, 16(1), 17-32.

- Hinds, J., and A. Joinson. 2019. "Human and Computer Personality Prediction from Digital Footprints." *Current Directions in Psychological Science* 28(2): 204–211. doi:10.1177/0963721419827849.
- Ienca, M., & Andorno, R. (2017). Towards new human rights in the age of neuroscience and neurotechnology. *Life sciences, society and policy*, 13(1), 1-27.
- Illes, J., & Sahakian, B. J. (Eds.). (2011). *Oxford handbook of neuroethics*. Oxford: Oxford University Press.
- Kiverstein, J. & Farina, M. (2012). Do sensory substitution devices extend the conscious mind? In F. Paglieri (Ed) *Consciousness in interaction: the role of the natural and social context in shaping consciousness*. John Benjamins Publishing. pp.19-40.
- Kiverstein, J. Farina, M. & Clark, A. (2015). Sensory Substitution. In M. Mohan (Ed) *The Oxford Handbook of Philosophy of Perception*. New York: Oxford University Press.
- Krueger, J. (2012). Seeing mind in action. *Phenomenology and the cognitive sciences*, 11(2), 149-173.
- Krueger, J., & Szanto, T. (2016). Extended emotions. *Philosophy Compass*, 11(12), 863-878.
- Lavazza, A. (2018). Freedom of thought and mental integrity: The moral requirements for any neural prosthesis. *Frontiers in Neuroscience*, 12, 82.
- Levy, N. (2007a). Rethinking neuroethics in the light of the extended mind thesis. *American Journal of Bioethics*, 7(9), 3–11.
- Levy, N. (2007b). *Neuroethics: Challenges for the 21st century*. Cambridge University Press.
- Lynch, M. (2013) *Brief of Michael P. Lynch as Amicus Curiae in Support of the Plaintiffs, American Civil Liberties Union v. James R. Clapper*. Available from: <https://www.aclu.org/legal-document/aclu-v-clapper-amicus-brief-michael-p-lynch-philosophy-professor-university> (Accessed, 19/09/2020)
- Locke, J. (1979). *Essay Concerning Human Knowledge*. Oxford University Press.
- Macpherson, F. (2018). Sensory substitution and augmentation: An Introduction. In *Sensory Substitution and Augmentation, Proceedings of the British Academy*, Macpherson, F. (ed.) Oxford: Oxford University Press.
- McCarthy-Jones, S. (2019). The autonomous mind: The right to freedom of thought in the twenty-first century. *Frontiers in Artificial Intelligence*, 19.
- McClelland, J., Rumelhart, D., & the PDP Research Group (Eds.). (1986). *Parallel distributed processing: explorations in the Microstructure of Cognition, Vols. I and II*. Cambridge: MIT Press/Bradford Press.
- Menary, R. (2007). *Cognitive integration: Mind and cognition unbounded*. Springer.

- Milojević, M. (2017). Extended personhood-rethinking property/person distinction. *Theoria, Beograd, 60*(4), 55-76.
- Palermos (2014). Loops, Constitution, and Cognitive Extension. *Cognitive Systems Research, 27*: 25-41.
- Palermos, S. O., (2018). Epistemic presentism. *Philosophical Psychology, 31*(3), 458-478.
- Peters, U. (2022). Reclaiming control: Extended mindreading and the tracking of digital footprints. *Social Epistemology*, 1-16.
- Reiner, P. B., & Nagel, S. K. (2017). Technologies of the extended mind: defining the issues. *Neuroethics: Anticipating the future*, 108-122.
- Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy, 101*, 389–428.
- Rupert, R. D. (2009). *Cognitive systems and the extended mind*. Oxford University Press.
- Roskies, Adina, "Neuroethics", The Stanford Encyclopedia of Philosophy (Spring 2021 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2021/entries/neuroethics/>>.
- Smart, P. R. (2012). The Web-Extended Mind. *Metaphilosophy, 43*(4), 446-463.
- Rowlands, M., & Mark, R. (1999). *The body in mind: Understanding cognitive processes*. Cambridge University Press.
- Rowlands, M. (2009). Extended cognition and the mark of the cognitive. *Philosophical Psychology, 22*(1), 1–19.
- Smart, P. Robert Clowes and Richard Heersmink (2017), “Minds Online: The Interface between Web Science, Cognitive Science and the Philosophy of Mind”, *Foundations and Trends in Web Science: Vol. 6, No. 1-2*, pp 1–232. DOI: 10.1561/18000000026
- Savulescu, J., & Bostrom, N. (Eds.). (2009). *Enhancement of human beings*. Oxford: Oxford University Press.
- Toon, A. (2014). Empiricism for cyborgs. *Philosophical Issues, 24*(1), 409-425.
- Weiskopf, D. A. (2010). The Goldilocks problem and extended cognition. *Cognitive Systems Research, 11*(4), 313-323.
- Wilson, R. & Lenart, B. (2014). Extended mind and identity. In J. Clausen & N. Levy (Eds.), *Handbook of neuro-ethics* (pp. 423–439). Dordrecht: Springer.
- Wittes, B. and Chong, J. (September, 2014). Our Cyborg Future: Law and Policy Implications <https://www.brookings.edu/research/our-cyborg-future-law-and-policy-implications/#footnote-4> [Accessed, 28 October, 2022].

