Extended knowledge overextended?*

Nikolaj Jang Lee Linding Pedersen

Jens Christian Bjerring

1. Mind and knowledge extended through technology?

In July 2020, Neuralink—the Elon Musk-led neurotech enterprise—received approval from the Food and Drug Administration for "breakthrough device testing" of their most recent brain-machine interface (BMI). On August 28, in a livestreamed update, Musk showed and described a coin-sized BMI that connects to the brain through approximately a thousand electrodes and is implanted at the top of the skull through a robot-driven surgical procedure.[1] While the initial commercial focus of Neuralink is on medical applications, the goal of human cognitive enhancement strongly guides the company's research and efforts. Musk is widely known for arguing that the only way for humans to compete with artificial intelligence is to seek cognitive enhancement through technology. To avoid becoming the 'pets' of super AIs, we must become cyborgs.[2]

In some sense, Musk claims, we are already cyborgs.[3] Technology may still be external to us but our reliance on smartphones and other computer technologyies is so heavy and regular that our computer devices act as if they are part of us. One of the key objetives of Neuralink is to *internalize* this kind of reliance through incorporation of technology into the human organism, effectively creating a symbiosis with artificial intelligence.[4] Indeed, during the August 2020 livestream, Musk confidently stated that "[i]n the future you will be able to save and replay memories"—suggesting an integration of a neurological device with storing and retrieval capacities into the human organism. But the ultimate goal is much more ambitious: it is to neurologically integrate into the human organism a minitiature computer that is able to perform the full range of tasks currently performed for us by external computer devices. The kind of neurological device envisioned by Musk is what Michael P. Lynch calls "neuromedia":

> Imagine you had the functions of your smartphone miniaturized to a cellular level and accessible by your neural network. That is, imagine that you had the ability to download and upload information to the internet merely by thinking about it, and that you could access, similar to the way you access information by memory, all the various information you now normally store on your phone: photos you've taken, contact names and addresses, and of course all the information on Wikipedia. (Lynch 2014: 299)

Setting aside issues about the technical construction and feasibility of such devices, neuromedia—and other sophisticated computer technologies—are philosophically significant.

* We are grateful to Davide Fassio, Jie Gao, Walter Hopp, Masashi Kasaki, Kevin Lynch, and Søren Overgaard for comments. We are likewise gratetul to Karyn Lai for her great work as the editor of this collection. A special thanks goes to Mog Stapleton for providing feedback on an earlier version of the paper. Her feedback prompted a substantial expansion of the introduction and considerable trimming of the other sections.

[1] For technical background, see Musk (2019)—a whitepaper on Neuralink. The full one-hour session is available at the Neuralink YouTube channel (https://youtu.be/DVvmgjBL74w).

[2] Interview at the Code Conference in 2016 (https://youtu.be/wsixsRI-Sz4 ; 56:50 onwards).

[3] Interview at the Code Conference in 2016 (https://youtu.be/wsixsRI-Sz4; 58:50 onwards).

[4] Interview at the Code Conference in 2016; launch presentation livestreamed on Neuralink's YouTube channel on July 16, 2019 (https://youtu.be/r-vbh3t7WVI).

First of, it is undeniable that modern technologies have vastly increased the number of true beliefs that people can quickly and reliably form. Take the pocket calculator as an example: by pushing a few buttons people can form a wide range of true beliefs about mathematics. Or take laptops or smartphones: by swiping a few times, a wide range of online resources such as *Nature, Science, Reuters*, *National Geographic,* and *Encyclopedia Britannica* offer people quick access to swaths of reliable information about the world. But according to an increasing number of philosophers, computer technology has a much more transformative effect.

Andy Clark and David Chalmers are famous—notorious even—for their defence of the so-called *extended mind thesis*. The basic idea is that our stock of beliefs can be extended through resources that are not among our biological onboard capacities. To get an intuitive feel for this idea, consider

> […] a normal case of belief embedded in memory. Inga hears from a friend that there is an exhibition at the Museum of Modern Art, and decides to go see it. She thinks for a moment and recalls that the museum is on 53rd Street, so she walks to 53rd Street and goes into the museum. It seems clear that Inga believes that the museum is on 53rd Street, and that she believed this even before she consulted her memory. It was not previously an occurrent belief, but then neither are most of our beliefs. The belief was sitting somewhere in memory, waiting to be accessed.
>
> Now consider Otto. Otto suffers from Alzheimer's disease, and like many Alzheimer's patients, he relies on information in the environment to help structure his life. Otto carries a notebook around with him everywhere he goes. When he learns new information, he writes it down. When he needs some old information, he looks it up. For Otto, his notebook plays the role usually played by a biological memory. Today Otto hears about the exhibition at the Museum of Modern Art, and decides to go see it. He consults the notebook, which says that the museum is on 53rd Street, so he walks to 53rd Street and goes into the museum.
> (Clark and Chalmers 1998: 12-13).

Prior to consulting her memory, it is uncontroversial that Inga has a standing or non-occurrent belief about the museum's location. What is not so uncontroversial is Clark and Chalmers' claim that Otto likewise has a standing belief about the museum's location—even before he consults his notebook. However, Clark and Chalmers argue that there is a parity between the information in Inga's biological memory and the information in Otto's notebook: both can easily and readily access the information as and when needed, and neither hestitates typically to act on the information when assessed. Given this parity, since we do not hesitate to attribute a non-occurrent belief to Inga about the museum's location, we should not hesitate to make a parallel attribtion in the case of Otto. Yet, if Otto has the same standing belief about the museum's location as Inga, the mind extends into the world: for Otto's beliefs sits in the notebook rather than in his biological memory.

There is an obvious generalization to computer technologies. Information stored online, on smartphones, or on laptops seems to be on par with information stored in biological memory. It, too, is easily and readily available as and when needed, and typically people do not hesitate to act on the information retrived from such external resources. You desire a great meal. You quickly and easily find the location of one the restaurants in Your Places on Google Maps and catch the subway in the right direction. You want to chat to an old college friend. You quickly and easily retrieve her number from the contacts on your phone and call her. Given this parity with information retrivial from biological memory, since we take biological memory to support non-occurrent beliefs, we should likewise take smartphones and laptops to support non-occurrent beliefs. In this sense, computer techonolgies can be part of the extended mind, or

computer technologies can extend cognition.[5]

As such, there seem to be three kinds of cognitive extensions:

(CE1) *Cognitive extension 1:*
Beliefs can be realized by onboard capacities as well as non-computerized external resources—so, e.g., by biological memory as well as notebooks.

(CE2) *Cognitive extension 2:*
Beliefs can be realized not only by non-computerized external resources but likewise by external computer devices—so, e.g., notebooks as well as smartphones and laptops.

(CE3) *Cognitive extension 3:*
Beliefs can be realized not only by external computer devices but likewise by computer devices integrated into the human organism—so, e.g., by smartphones and laptops as well as neuromedia.

Each type of cognitive extension is significant in relation to *the nature of belief*. If we accept that beliefs—and other epistemic notions such as knowledge—extend beyond the skin and skull of cognizers, we also accept a bifurcation in the nature of belief states. Whichever way cognitive extension may be achieved—whether through a notebook, a smartphone, or a neuromedium—there is a fundamental *qualitative* difference between non-extended and extended beliefs. While the former is realized through information stored by onboard biological capacities, the latter is realized through information on external non-computerized resources or computer devices.

However, while all of (CE1), (CE2), and (CE3) in some sense involve acknowledging a new *kind* of belief, it seems that only (CE2) and (CE3) involve granting individuals belief on a much larger *scale*. In principle, we could always carry with us notebooks or other non-computerized resources in which we mechanically record information in ways that make it easily and readily available as and when needed. But there is quite a low upper bound on how much information we can record, store, and retrieve in this way. Recording, storing, and retrieving information only become fast, convenient, and feasible on a grand scale when computer devices are involved. It is thus only with (CE2) and (CE3) that extended beliefs make a considerable *quantitative difference*.

Compared to pre-internet times or the early days of the commercial internet (56K dial-up modems, anyone?), it is now much easier to look up information. We have swaths of information right at our fingertips at any given time. Three decades ago, if you wanted to learn when and how Napoleon died, you might well have had to make your way to the library, pull a book down from the shelves, and look it up. Today, the answer to many such questions are just a few swipes away. Suppose someone accesses true information about Napoleon through a reliable online source on her smartphone and stores that information on the phone. Granted the extended mind thesis, once the information is no longer occurrently believed, the person still qualifies as non-occurrently believing this information—even if the person is unable to recall the information by exercising onboard capacities. Why? Because, just like information in biological memory, information stored on the person's smartphone is easily and readily available as and when needed.

More generally, then, the extended mind thesis supports the striking idea that computerized memory can be your memory. And since computerized memory can support

[5] See also Chalmers (2008), the foreword to Clark's *Supersizing the Mind* (2008). Chalmers' 2011 TEDxSydney talk was entitled "Is Your Phone Part of Your Mind?" and gives the same argument (https://youtu.be/ksasPjrYFTg). Musk has said repeteadly in interviews that we are already cyborgs—see, e.g., the interview at the Code Conference 2016 (https://youtu.be/wsixsRI-Sz4; 58:50 onwards).

belief in the same way that biological memory supports belief, the extended mind thesis supports the striking idea that any person with ready and reliable access to a connected computer device is in a position easily to acquire and retain a vast amount of true beliefs about a wide range of subject matters.[6] This is why devices such as smartphones and laptops carry great quantitative significance for the number of beliefs that individuals have.[7]

We might worry that this quantitative implication of extended belief causes an unwarranted *cognitive bloat*—as does, indeed, Åsa Wikforss:

> If I carry my iPhone [and it's easily and readily available as and when needed] … and I have downloaded (and accepted) everything there is to know about the Icelandic fishing industry, does it really follow that this information constitutes my standing beliefs? For instance, does it really follow that I have the standing belief that in 2011 the export production of marine products amounted to ISK 252 billion and increased in value by 14.4% from previous year?
>
> Once we take the step from extended belief to extended knowledge the worry about 'bloat' becomes more acute. If the downloaded information belongs with my standing beliefs does it also follow that I now *know* everything there is to know about the Icelandic fishing industry? Have I suddenly become an expert in the topic? (Wikforss 2014: 465.)

Katalina Farkas expresses the same kind of concern[8]:

> … if one never takes a decision without consulting their Filofax, the Filofax carries beliefs and intentions; if one is "unusually computer-reliant, facile with the technology, and trusting", parts of the Internet may carry beliefs; in the case of an unusually interdependent couple, the beliefs of one's spouse may act as standing states of the other
>
> Now the worry is that by going this far, we'll be forced to let in *other* cases that create tensions in our normal notion of a subject or a self. Suppose someone, call her Lotte, always carries an electronic reading device. She downloads a 37 volume history of Europe with a quick search function, from a source she trusts completely. If any question about the history of England comes up, Lotte consults the book … Does Lotte thereby acquire all the beliefs, and hence does she become an expert on the history of England? To say that she does is at least prima facie problematic. We do want to distinguish between an impressively erudite scholar of English history who mastered the subject, and Lotte, who simply stores a file on a device. It seems the learning does truly belong to the historian, but not really to Lotte. (Farkas (2012: 444))

But why is cognitive bloat troublesome?

First of all, cognitive bloat seems counter-intuitive: is it really possible to become an expert on some complicated topic like the Icelandic Fishing Industry or European History merely in virtue of being appropriately related to certain technologies? As Wikforss and Farkas put it:

---

[6] In Bjerring and Pedersen (2014), we talked about this consequenece of the extended mind thesis in terms of *restriced omniscience:* a person enjoys restricted omniscience with respect to a particular, fairly specific subject matter—say, past Academy Award winners or NBA standings and statistics—if the person has complete (or close to complete) extended knowledge with respect to that subject matter.

[7] Put differently, computer devices can result in a significant upgrade to your biological memory. Earlier we spoke of boosting memory through neuromedia, and here we say that computer devices, whether external or internal, can result in a memory upgrade. It is important to qualify this. We are *not* saying that human memory is small. It is widely agreed that it is *huge*—about 2.5 petabytes (or a million gigabytes). The totality of our memories takes up much, much more space than what is available on smartphones and other computer devices we have access to. There are many types of memories and different ways that they are stored. What we have in mind when making the upgrade-related claims is quite specific: think of how much text you can store on a regular laptop or smartphone. Now consider whether you could store *all* of that information and retrieve it from biological memory. Probably not. Computer memory involves an upgrade in the sense that, counting it as yours, it puts you in a position to *recall* many things that you would otherwise not be able to recall.

[8] See also Rubert (2004, Sect. IV) and Sprevak (2009, Sect. V).

> Of course, it is part and parcel of the extended mind thesis that belief and knowledge does extend in rather surprising ways. However, no one would take it to be a virtue of the theory that the extension is unchecked, clashing with widely shared intuitions about what does and what does not count as a subject knowing that *p*. (Wikforss 2014: 461.)

> To say that [Lotte knows everything in the 37 volumes of history downloaded] is at least prima facie problematic. We do want to distinguish between an impressively erudite scholar of English history who mastered the subject, and Lotte, who simply stores a file on a device. It seems the learning does truly belong to the historian, but not really to Lotte. (Farkas (2012: 444))

The worry here is that allowing technologies to extend beliefs makes it *too easy* to attain knowledge and status as an expert. That is, tech-driven cognitive extension extends our knowledge greatly beyond ordinary intuitions and so obliterates the distinction between experts and novices. As reflected by the passages from both Wikforss and Farkas, cognitive bloat puts pressure on a natural way of thinking about expertise and knowledge: an expert concerning a given subject *X* is someone who has a lot of knowledge about *X*. Thus, an expert concerning organic chemistry is someone who knows a lot about organic chemistry, and an expert concerning the history of England is someone who knows a lot about that particular subject. However, can one go from novice—or even *ignoramus*—to expert by downloading or storing information on a smartphone? The transition from novice to expert would seem to require more than that. Nonetheless, if expertise is thought of *merely* in terms of knowledge, it would seem that computer-driven cognitive extension puts expertise with respect to a very wide range of subject matters within easy reach of ordinary subjects. This seems unpalatable to some.

Second, there is a related conceptual worry that we might have to say that

> … human cognition extends in certain cases and to a certain degree, but they usually want to avoid claiming too much cognitive extension. Too much extension might make our existing mental concepts pointless, absurd, or otherwise unfit for purpose. We might end up saying you believe everything on the Internet … (Sprevak 2019, Sect. 3)

Intuitively, all truths on reliable online platforms are not within our epistemic reach. If all such truths qualified as beliefs—or, stronger yet, knowledge—this might be thought to render the concepts of belief and knowledge absurd or pointless. Indeed, one might think that this is a *reducio* against the very idea of extended belief and knowledge.[9]

Third, we might worry that the accommodation of extended beliefs radically changes the nature of epistemology. For if intutions reflect our ordinary concept of belief, and if extended beliefs are counter-intuitive—in part due to worries about cognitive bloat—then cognitive extension militates against our ordinary epistemological concepts. In this case, one might hold that extended belief calls for a *revision* of our ordinary concepts. Instead of revision, one might alternatively insist that the concept of extended belief is a *new* epistemological concept. As a consequenece, extended epistemology has a subject matter that is different from epistemology as traditionally understood. Either way it seems that the extended mind thesis radically changes the nature of epistemology. Either epistemology concerns a *revised* concept of belief or it concerns a *new* concept of belief, and, accordingly, it divides into two sub-disciplines concerned with these different concepts.

---

[9] Interestingly, this complaint suggests the existence of an intimate connection between the qualitative and quantitative significance of cognitive extension: non-onboard capacities and resources such as notebooks and external computer devices do not support extended belief because it would imply massive cognitive bloat. A point about the quantity or scale of cognitive extension supports a substantial qualitative conclusion.

In the remaineder of this paper, we will discuss the challenge of cognitive bloat in detail. After stagesetting in section 2, we consider in section 3 Wikforss' strategy for resisting cognitive bloat. While Wikforss' strategy might be effective in blocking cognitive bloat from (CE2) types of extension, it does not nothing, we argue in section 4, to block cognitive bloat from (CE3) types of extension. In section 5, we turn to J. Adam Carter and Jesper Kallestrup's cluster-model functionalism—a view which, to a considerable extent, is designed to deal with issues of cognitive bloat. In section 6, in addition to raising several issues regarding cluster-model functionalism, we argue that cognitive bloat from (CE3) types of extension is not ruled out by their framework. So while it remains an open question to us whether cognitive bloat is indeed a problem for the extended mind thesis or not, we conclude that neither the framework proposed by Wikforss nor that proposed by Carter and Kallestrup avoid cognitive bloat. In fact, in light of (CE3) types of cognitive extensions, we are doubtful that cognitive bloat can be avoided in the functionalist landscape of the extended mind thesis. For some this will be enough to abandon the extended mind thesis. For others it will be an interesting observation about our future lives as cyborgs.

2. From extended mind to extended knowledge and restricted omniscience

In this section we introduce the idea of extended mind in more detail and then move on to extended knowledge and cognitive bloat.

The Inga and Otto case introduced earlier is the key example of extended mind in Clark and Chalmers (1998). Their commitment to extended mind is driven by an endorsement of *functionalism* about belief: beliefs are characterized in terms of their functional role. In particular, beliefs combine with desires to explain action. What explains Inga's going to the MoMA? Her desire to do so together with the belief that the MoMA is on 53rd Street. Otto carries out the same action as Inga: he goes to 53rd Street. Like Inga he has the desire to go to MoMa. The information in Otto's notebook seems to play the same role as Inga's belief that MoMa is on 53rd Street: it combines with a desire to go to MoMa to explain the action of going to 53rd Street. To further support the idea of sameness of functional role, consider a scenario in which the belief about MoMA's relocation was removed from Inga's biological memory. In that case, she would not have gone to 53rd Street. Similarly, if the address of the MoMA were to be deleted from Otto's notebook, he would not go to 53rd Street either.

Given functionalism about belief, it seems that, by parity of functional role, the information in Otto's notebook sustains a non-occurrent belief concerning MoMA's location if the information in Inga's biological memory does. Yet, Otto's belief depends constitutively on something in the external environment—namely, the notebook. In this sense, Otto's belief is extended, and in this sense, the mind extends into the world. Information stored in external devices qualifies as non-occurrent beliefs since, functionally, it plays the same role as information stored in biological memory.

It is natural to wonder whether any external resource whatsoever can sustain extended beliefs. Clark and Chalmers do not think so. They discuss the following four conditions as constraints on extended belief:

(C_1) *Typical invocation*: the resource is readily available and typically invoked.

(C_2) *Automatic endorsement*: any information retrieved from the resource is more or less automatically endorsed.

(C_3) *Easy access*: the information contained in the resource is easily accessible as and

when required.

(C_4) *Past endorsement*: the information in the external resource has been consciously endorsed by the subject at some point in the past and is contained in the resource as a result of this endorsement.

The first three conditions are often called the 'trust and glue conditions' since they are meant to ensure that the external source is trusted by the subject and 'glued' to the subject.[10]

How does Otto's notebook fare with respect to C_1-C_4? All conditions appear to be satisfied. Otto always carries the notebook and relies on it for navigating in his daily life and environment. Hence, the notebook is readily available and typically invoked. So C_1 is satisfied. Otto automatically endorses information retrieved from the notebook. He treats it as trustworthy and does not subject retrieved information to critical scrutiny. This is reflected by him not hesistating to act on information he retrieves from the notebook. So C_2 is satisfied. Since Otto always carries the notebook with him, the information in it is easily accessible as and when needed. So C_3 is satisfied. C_4 is likewise satisfied since all information in the notebook has been consciously endorsed in the past and is contained in the notebook as a result of this endorsement.[11]

Clark and Chalmers take C_1-C_4 to be jointly sufficient for an external resource to support extended beliefs, while they regard C_1-C_3 as individually necessary for an external resource to sustain extended beliefs. However, since they take seriously the idea that C_1-C_3 might be jointly sufficiently for extended belief, they also take seriously the idea that C_4 might not be a necessary condition. They cite the possibility of subliminal perception or memory tampering as a reason to refrain from regarding C_4 as a necessary condition. For current purposes, we proceed on the assumption that conditions C_1-C_4 are individually necessary and jointly sufficient for an external resource to sustain extended beliefs. Including C_4 clearly makes it more demanding for an external resource to support extended beliefs. So by including C_4 we make our main conclusions stronger.

Epistemologists have taken a keen interest in Clark and Chalmers' work, using it as a platform to launch the idea of *extended knowledge*. Some epistemologists have argued that knowledge is extended by appeal to the idea of *extended cognition* (Clark and Chalmers (1998)). The idea is simple: if a knowledge-generating cognitive process is partly driven by something in the external environment, then knowledge is extended. The extendedness of knowledge derives from the process through which the belief is generated or sustained. While part of what drives cognition is extended, the state of belief is not.[12]

Other epistemologists explore a diferent path to extended knowledge: through extended mind. The basic idea is that the extendedness of knowledge derives from the extendedness of belief, i.e. the state that qualifies as knowledge. In Bjerring and Pedersen (2014), we explored this path by arguing that the combination of the extended mind thesis and reliabilism carries a

[10] Clark and Chalmers (1998: 17). Our formulation of the three first conditions follows Clark (2010: 46). The label "trust and glue" is due to Clark.

[11] Clark and Chalmers (1998: 17).

[12] See, e.g., Kelp (2013), Pritchard (2010), Palermos (2011, 2014, 2015). As for examples of extended cognition, think of the game Tetris where buttons are used to rotate figures in order to determine whether they fit into empty slots in a structure, or the task of determining the product 13×17 by using a calculator. According to Clark and Chalmers, the following *parity principle* forces a non-discriminatory approach to internally and externally driven cognition: "If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process." (Clark and Chalmers 1998: 8) Since we would not hesitate to take mental rotation and mental calculation as part of the cognitive process, we should not hesitate to recognize the external devices as part of the cognitive process either. Cognitive processes thus extend into the external environment.

commitment to extended knowledge. Based on this argument for extended knowledge, we argued that ordinary subjects can easily obtain various forms of *restricted omniscience*. That is, ordinary subjects can easily come to know all truths about some fairly specific subject matter—e.g., the names and birthdates of all past US presidents. Let us briefly rehearse the argument here.

Reliabilism is a widely held view in epistemology, which we, for present purposes, can characterize as follows:

(R) *Reliabilism:* A subject *S* knows that *p* if
- (i) *S* believes that *p*;
- (ii) *p* is true; and
- (iii) *S*'s belief that *p* is formed or sustained by a reliable process.

Note that conditions (i)-(iii) are given as jointly sufficient for knowledge. A reliabilist would also take each condition to be individually necessary. However, as shall transpire below, what is relevant to our argument is the sufficiency claim. Hence (R) will do for the purposes of our argument.[13]

With the extended mind thesis and reliabilism on the table, the argument for extended knowledge is straightforward:[14]

Extended Knowledge Argument:

(1) If subject *S* has a belief that *p*—whether extended or not—that is formed or sustained by a reliable process and is true, then *S* knows that *p*.
(2) *S* has a true, extended belief that *p* that is sustained by a reliable process.
(3) Therefore, *S* knows that *p*.

The first premise is an epistemological thesis: reliabilism about knowledge. The second premise claims that there is a belief that satisfies the three conditions identified as being sufficient for knowledge. The second premise also makes it clear why the path to extended knowledge goes via extended mind: it expclitly states that the relevant belief is an extended belief, i.e. a belief sustained by an external resource. So, e.g., assuming that Otto mostly gathers information from reliable sources, true information in his notebook qualifies as knowledge because sustaining beliefs via the notebook is a reliable belief-sustaining process.

As formulated, the Extended Knowledge Argument leaves an important question open: what is the *quantitative significance* of extended knowledge? In Bjerring and Pedersen (2014), we argued for the following thesis:[15]

> For any subject *S*, *S* knows each true piece of information contained in a reliable external resource that satisfies conditions C_1 to C_4 relative to *S*.

This is the juncture at which computer devices become highly significant. For given such devices, it is clear that ordinary subjects are in a position easily to acquire a wealth of extended knowledge—that, indeed, ordinary subjects can easily come to enjoy various various forms of restricted omniscience:

---

[13] Classic works on reliabilism include Goldman (1979), (1986). It is standard also to include a no-defeaters condition. We leave it out here for ease of exposition. However, nothing hangs on this simplification, as the arguments and considerations satisfy this type of condition or can be straightforwardly modified to do so.
[14] Bjerring and Pedersen (2014: 27).
[15] Bjerring and Pedersen (2014: 30).

(RO) *Restricted omniscience:*
Given extended belief and knowledge, any person with ready and reliable access to a connected computer device is in a position easily to acquire and retain restricted omniscience about a wide range of subject matters.

The following case helps to illustrate this point:[16]

Lone:
Lone spends a month cutting and pasting information from various online resources into a document that she keeps on her smartphone. Lone cuts and pastes information from online resources that she takes to be reliable and which are, indeed, reliable (i.e. they contain a high proportion of true pieces of information). These resources include encyclopedia, news reports, and scientific journals, among others. As she cuts and pastes, Lone consciously endorses each piece of information. After a month of cutting and pasting, Lone has cut and pasted a vast amount of information—including, but also greatly exceeding, the list of all past Super Bowl winners; all past NFL, NBL, NBA, and NHL winners; all birthdates of past Academy Award winners, British prime ministers, and French presidents; the first 10,000 entries in the multiplication table of 17; and the daily conversion rate from US dollars to Korean won for the past 20 years. All of this information—and much, much more—is included in the document on Lone's phone. The document is easily searchable, and Lone has easy and ready access to the document, has endorsed each piece of information that it contains, and indeed, each piece of information is there as a result of this endorsement. Lone frequently relies on the document in her everyday life, automatically endorses the information that she retrieves from it and does not hesitate to act on it. Since Lone herself has created the information-storing document and has also herself cut and pasted the information it contains, she is aware of the information-storing source and the sources from which information has been collected. Lone takes her way of collecting, storing, and retrieving information to be reliable. Since the sources she used are reliable and since information cut and pasted is preserved and retrieved in the form it was entered into the document, Lone's way of collecting, storing, and retrieving information is indeed reliable.

By design the document on Lone's phone satisfies conditions C_1-C_4. The information cut and pasted into the document thus qualifies as extended beliefs. Likewise, by design, the document lcontains mostly true information, as she cuts and pastes from reliable online resources. By the Extended Knowledge Argument, each true piece of information in Lone's document qualifies as extended knowledge. Lone thus enjoys various forms of restricted omniscience. She knows everything about some specific subject matter—including, e.g., the four biggest commercial sports in the US: football, baseball, basketball, and hockey.

Now, there is nothing extraordinary about Lone. She is very much like any ordinary person who owns a constantly connected smartphone. The punchline—not surprisingly, perhaps—is that ordinary people are in a position to easily attain various forms of restricted omniscience. Let us emphasize that most people have not *actually* spent a considerable amount of time cutting and pasting information into a document on their laptop. So most people are not *actually* restrictedly omniscient in exactly way that Lone is. However, if our argument is sound, it is possible for ordinary subjects easily to attain various forms of restricted omniscience. The Lone

[16] The Lone case is like the case called "Cut and Paste" in Bjerring and Pedersen (2014: 29), except that it makes explicit the subject's attitude towards the existence and reliability of the information-sustaining source. This added layer of explicitness will the helpful for the purposes of our discussion in Sects. 5 and 6.

case, after all, specifies a straightforwardly executable procedure for doing so. Furthermore, if our argument is sound, it *does* seem that large bodies of true information stored externally *actually* qualifies as knowledge—and so, that ordinary subjects actually know much more than we ordinarily take ourselves to know. Just think of all the true information that you have stored on your phone—in your contacts, Dropbox, Google Drive, or what have you—which has been retrieved from reliable sources and which is readily and easily available as and when needed.

As stated above, our argument incorporates the assumption that condition C_4 must be satisfied in order for information in an external resource to qualify as belief. If, however, C_1-C_3 were jointly sufficient on their own, it would have a significant impact on the scope of extended knowledge. For without the requirement of past endorsement, extended beliefs would be a lot easier to come by, *and* many of these beliefs would satisfy premise (2) of the Extended Knowledge Argument—that is, they would be true and sustained by a reliable process. Why? Because many ordinary subjects enjoy access to a wide range of reliable online resources such as *Nature*, *Science*, *Reuters*, *National Geographic*, and *Encyclopedia Britannica* through their mobile phone. They enjoy such access in a way that satisfies C_1-C_3: the resources are typically invoked or relied on for certain kinds of information, they are easily and readily available as and when needed, and information retrieved from them is automatically endorsed. Since the online resources in question contain a lot of information, they sustain a great number of extended beliefs. And since they are *reliable* online resources, these extended beliefs are reliably sustained. Given reliabilism about knowledge, this means that any true piece of information contained in these resources will qualify not only as belief but as knowledge. Many ordinary subjects thus possess a wealth of knowledge concerning science, current affairs, history, and geography. Indeed, they possess much more than that. For, again, true information available through any realible online resource satisfying C_1-C_3 would qualify as knowledge. So count in any website that reliably stores records of the weather, exchange rates, standings of soccer leagues around the world, and what have you.

The possibility of restricted omniscience makes vivid the counterintuitive feel of cognitive bloat: it makes vast amounts of knowledge *too cheap*, *too easy*, and *too quick* to attain. So how might one address concerns about cognitive bloat or restricted omniscience?[17]

3. Against cognitive bloat (I): no belief

One strategy for dealing with cognitive bloat is to argue that the relevant external devices fail to sustain beliefs. Since belief is required for knowledge, the relevant external devices thus fail to deliver knowledge. Wikforss (2014) pursues this kind of strategy.

But why deny that external resources like those used by Otto and Lone sustain beliefs? Because they do not satisfy the following condition, which, according to Wikforss, is a requirement for belief proper:[18]

> (C_W) *Evidence-sensitivity and direct interaction*: the information in the external resource is appropriately sensitive to evidence and interacts directly with the subject's internally stored beliefs and desires.

[17] Other discussions of cognitive bloat or intimately related issues include Allen-Hermanson (2013), Carter and Kallestrup (2020), Farkas (2012), Ludwig (2015), Lynch (2014, 2016), Pritchard (2010, 2018a, 2018b, 2018c), Rupert (2012), Smart (2012), Sprevak (2009, 2019), Wikforss (2014). More on Wikforss (2014) and Carter and Kallestrup (2020) below.

[18] Wikforss (2014, Sect. 3).

There are thus two ways in which information in external resources can fail to satisfy C_W. The information might not be appropriately sensitive to evidence, or it might not be appropriately connected to beliefs, desires, and actions to fulfill the functional role of belief.

Let us consider sensitivity to evidence first. It is a widely held view that the formation, destruction, and sustainment of belief are regulated by evidence.[19] For instance, if someone receives evidence that a given belief *p* of hers is false, the person's belief system gets updated: *p* is discarded and other adjustments are made automatically in light of the evidential, logical, and probabilistic relations that exist between the belief that *p* and other beliefs. Suppose Sophie, a middle school kid, holds the following beliefs:

p_1 All birds can fly.
p_2 Ostriches, a certain bird species, can fly.

Now suppose that during biology class Sophie learns that ostriches, a certain bird species, cannot fly. In that case, Sophie discards her belief in p_2, and, given its logical relation to p_2, p_1 is automatically discarded as well.

This kind of updating does not happen automatically in the case of Otto. If p_1 and p_2 were entries in the notebook, then, upon learning $\sim p_2$, p_1 would not automatically be discarded and replaced by $\sim p_1$. In order for entries in the notebook to be updated, Otto needs to take action. Upon learning $\sim p_2$, that is, Otto would have to manually enter $\sim p_1$ into the notebook in order to repect the new evidence. This process seems to be all but automatic. By the lights of Wikforss, the information in Otto's notebook thus fails to be appropriately sensitive to evidence. Hence it cannot qualify as belief. The same point seems to apply to information in Lone's document.

Next consider the functional role of belief, as understood in terms of its relation to action and desire. Wikforss grants Clark and Chalmers their functionalist starting point, conceding in particular that beliefs in conjunction with desires play an essential role in explaining action.[20] However, she argues that the information in Otto's notebook does not play this role. This is because information in the notebook does not interact *directly* with Otto's internally stored beliefs and desires:[21]

> In the case of Inga there is a very direct interaction between her beliefs and desires. Thus, after reading about the exhibit at MoMA, a simple piece of practical reasoning leads to her action:
>
> Desire 1: *I go to MoMA.*
> Belief 1: *MoMA is on 53rd Street.*
> Action: *I go to 53rd Street.*
>
> In Otto's case, by contrast, the connection between the information in the notebook and his internally stored beliefs and desires is mediated by a complex set of other beliefs and desires, including second-order beliefs. After reading about the exhibit at MoMA Otto reasons along the following lines:
>
> Desire 1: *I go to MoMA.*
> Belief 1: *I do not remember where MoMA is.*
> Desire 2: *I find out the location of MoMA.*
> Belief 2: *This kind of information is written down in the notebook.*
> Desire 3: *I look in the notebook.*
> Belief 4: *It says in the notebook that MoMA is on 53rd Street.*
> Belief 5: *I believe everything that's written in the notebook.*
> Belief 6: *MoMA is on 53rd Street.*

[19] See, e.g., Shah and Velleman (2005).
[20] Wikforss (2014: 469).
[21] Wikforss (2014: 470-471).

Action: *I go to 53rd Street.*

The structure of Otto and Inga's respective lines of practical reasoning seems to be quite different. Inga has a standing attitude—*belief*—towards the information stored in her biological memory, *viz*. that MoMA is on 53rd Street. This belief becomes occurrent and, given her desire to go to MoMA, Inga acts on it. The standing belief and the occurrent belief share their content and the transition from one to the other is a matter of immediate recall, not a mediated transition that goes through second-order beliefs *about* Inga's memory. For Otto, however, the path to the belief that MoMA is on 53rd Street is *mediated*. He arrives at this belief through several second-order beliefs: beliefs *about* the notebook containing the relevant kind of information, and beliefs about him, Otto, believing everything in the notebook. Similarly, the information in Otto's notebook is *just information*. Otto has no standing attitude towards it, according to Wikforss, because it takes several second-order beliefs about the notebook for that information to feature in an occurrent belief that can combine with his desire and result in him going to 53rd Street. Contra Clark and Chalmers, Wikforss thus maintains that the information in the notebook cannot play the functional role of belief.[22]

So according to Wikforss, the information in Otto's notebook fails to satisfy the additional condition C_W on belief. As such, Otto's notebook does not sustain beliefs. The same point applies to the information contained in Lone's document, in documents on Dropbox or on mobile phones—indeed, the same point applies to any external device or resource that fails to satisfy C_W. Wikforss thus blocks concerns about restricted omniscience or cognitive bloat at the level of belief: there is no explosion or overextension of knowledge quite simply because knowledge requires belief *and* standard information-storing external devices fail to sustain belief.[23]

4. Against: against cognitive bloat (I)

Ultimately, we do not find Wikforss' response to the issue of cognitive bloat effective. For imposing C_W as an additional condition on belief does not amount to a principled rejection of cognitive bloat. The effectiveness of Wikforss' argument is contingent on the current state of technological development. Even if sound, the argument only shows that external computer devices do not *currently* sustain beliefs. It does not show that computerized devices *cannot* do so.

On our way of carving up the landscape, Wikforss' condition C_W might prevent (CE1) and (CE2) type cognitive extensions from bloating, but it does nothing to address (CE3) type extensions like those made possible by neuromedia. Interestingly, Wikforss seems to agree:

> … once we accept a functionalist account of belief there is a principled possibility of external vehicles. However, the functional demands are harder to meet than is usually assumed, and go well beyond the trust and glue conditions. What is required is that the external information plays the role of belief in ordinary folk psychological reason explanations, and this condition is not met simply by the subject's having easy and reliable access to the information. What is required in addition, is that the external information interacts directly with the subject's further beliefs and desires. It is possible that in the future a device could be built (perhaps replacing part of a damaged brain) which allows the externally stored information to fully interact with the internally stored information. However, something much more sophisticated

[22] The functional role of belief in explaining action together with desire relates to *practical reasoning*. As just seen, Wikforss argues that the information in Otto's notebook is unfit to feature in practical reasoning in the way that beliefs are supposed to do. She further argues that the information in the notebook is unfit to play the role of belief in *theoretical reasoning*. See Wikforss (2014: 472-474).

[23] Wikforss (2014: 475).

> would be needed than a simple notebook or an iPhone. (Wikforss (2014: 475))

So Wikforss herself suggests that a replacement device integrated into a damaged brain might satisfy C_W. The role of such a device is similar in certain respects to that of neuromedia. By design, neuromedia are intended to integrate seamlessly with our natural or onboard capacities. Lynch supposes this integration to take the form of thought-controlled retrieval and uploading of information to the internet and access to stored information in a way similar to recall from memory.

More precisely, the specific kind of neuromedia that Lynch envisions involves two kinds of integration. For internet-related tasks such as Googling or posting a message on your Facebook wall, the subject is still very much a *user* of a clearly separate resource, service, or platform. Neuromedia integrate these resources, services, and platforms by facilitating a thought-controlled link to them. In this sense, the link is internal as well as immediately and readily available. Still, this kind of integration goes with a phenomenology that is distinct from exercising natural or onboard capacities. For the subject will very much be aware that she is using a separate resource, service, or platform. But Lynch also imagines a second kind of integration: namely accessing stored information in a way similar to retrieving information from memory. In this case, the integration of neuromedia is seamless. The phenomenology of information-retrieval will be *similar* to the phenomenology of recalling information from biological memory.

Duncan Pritchard also discusses neuromedia:

> I want to explore one particular kind of technological development that might well be right around the corner: neuromedia … information-processing technology that is so seamlessly integrated with our on-board cognitive processes that the subject is often unable to distinguish between her use of those on-board processes and the technology itself. The relationship of the subject to the technology is consequently no longer one of subject to instrument but rather "feels" like a technological extension of her normal cognitive processes. (Pritchard 2018a: 328-329.)

Pritchard's neuromedia are slightly different from Lynch's. Given his focus on the second kind of seamless integration that Lynch has in mind, Pritchard uses the term "neuromedia" in a narrower, stronger sense than Lynch. This difference becomes apparent when Pritchard describes the kind of information search that, these days, is typically carried out using Google (or some other search engine). For Lynch this kind of search will be integrated with the subject in the sense that it is thought-controlled. Instead of typing "How many inhabitants does Berlin have?" into Google, you can initiate the search and receive the answer through thought processes alone. But the phenomenology is still one of the subject's *using* a clearly separate service, platform, or resource. Pritchard, on the other hand, envisions that accessing the relevant information and the phenomenology of doing so are "just like remembering".[24]

Both Lynch and Pritchard, however, take neuromedia to deliver knowledge. This means that they must grant that neuromedia results in belief, as belief is necessary for knowledge. We also take it that they would both grant the reliability of the relevant belief-generating or belief-sustaining process. As such, neuromedia can deliver extended knowledge in the sense that we have been discussing it here, i.e. as true belief sustained by a reliable process.[25]

---

[24] Pritchard (2018a: 329).

[25] Lynch uses the term "Google knowledge" for the kind of knowledge generated through neuromedia. He is keen to emphasize that we should be critical of such knowledge and recognize its limitations—in particular, that it fails to deliver understanding. For a book-length discussion of this issue, see Lynch (2016). Pritchard likewise argues that we need to recognize the limitations of knowledge generated through external devices or neuromedia. Among other things, even if they result in swaths of knowledge, this type of knowledge does not contribute towards achievement of the chief epistemological aim of education: to build intellectual character. This aim, according to

So the question is whether the kind of neuromedia that Lynch or Pritchard envision can deliver knowledge by Wikforss' lights? This boils down to the question whether they satisfy conditions C_1-C_4 and C_W. We think that they fail to do so because they fail to satisfy the condition of past endorsement C_4. Both Lynch and Pritchard assume that neuromedia will give us easy and ready access to swaths of information. This is clear from the fact that Lynch explicitly describes neuromedia as providing a thought-controlled link to the internet and giving memory-like access to all of Wikipedia. Similarly, Pritchard talks explicitly about neuromedia providing answers to questions that, these days, we would seek to answer by using search engines. The potential amount of information that neuromedia put right at our finger tips is thus so vast that no ordinary subject can have endorsed all of it in the past. So Lynch and Pritchard's neuromedia cannot in general satisfy C_4.

In light of this, let us restrict attention to neuromedia paired with storage devices that contain only information endorsed by the subject:

> *Lone*$^+$
> Lone$^+$ is knowingly cognitively enhanced through the implantation of a neurological device, and she is informed about how it works. The device includes a hard drive and a unit that allows Lone$^+$ to search the internet by thought alone and to store information retrieved on the hard drive. The hard drive does *not* come loaded with information. It only contains information that gets stored as a result of Lone$^+$'s activities. The phenomenology of storing and retrieving information from the hard drive is indistinguishable from storing and retrieving information from biological memory. Furthermore, the device builds in an algorithm that updates the hard drive on the basis of evidential, logical, and probabilistic relations between new information and the information on the hard drive. Lone$^+$ spends a month accessing information from various reliable online resources—including encyclopedia, news reports, and scientific journals—and she endorses their reliability. Upon accessing each piece of information, Lone$^+$ endorses it, and it gets stored on the neurologically implanted hard drive. After a month a vast amount of information has been accessed and stored. All of it now sits in the hard drive. Lone$^+$ has easy and ready access to the information in the hard drive and the information is there as a result of her past endorsement. She frequently relies on the hard drive in her daily life, automatically endorses the information she retrieves from it and does not hesitate to act on it. Since Lone$^+$ has knowingly been cognitively enhanced through the implantation of a neurological device, and since she herself has collected the information it contains by searching the internet (in thought), she is aware of the information-storing source and the sources from which the information has been collected. Since the sources Lone$^+$ used are reliable and stored information is retrieved in the form it was entered, Lone$^+$'s way of collecting, storing, and retrieving information is indeed reliable.

The Lone$^+$ case replaces the document in the Lone case with a neurologically integrated search and storage device. The acts of accessing information, endorsing it, and cutting and pasting it into the document have been replaced by the acts of accessing information, endorsing it, and its being stored in the hard drive in a way that is phenomenologically indistinguishable from when information gets stored in biological memory. Likewise, the acts of searching the document, locating relevant information, and (re-)endorsing it have been replaced by a process of retrieval that is phenomenologically indistinguishable from recalling information from

---

Pritchard, is a matter of cultivating and exercising onboard capacities. See Pritchard (2018a) for an extended argument to this effect; see likewise Pritchard (2017). Interesting as they may be, we have relegated these observations to a footnote because what matters for present purposes is the fact that both Lynch and Pritchard grant that neuromedia sustain belief.

biological memory.

While the neurologically integrated device in Lone$^+$ quite clearly satisfies conditions C_1-C_4, the question is whether it also satisfies Wikforss' condition C_W? Arguably it does. It seems to be appropriately sensitive to evidence and integrate or interact directly with beliefs and desires, as demanded by C_W.

First, the information stored on Lone$^+$'s neurologically integrated search and storage device device seems appropriately sensitive to evidence. By construction, the device updates automatically via its built-in updating algorithm. Suppose, e.g., that Lone$^+$ endorses the information that Jimi Hendrix was born on September 18 after consulting a music website (in thought, remember!) and forms an occurrent belief to that effect. Since she believes that Virgos are born between August 23 and September 22, Lone$^+$ reasons from the belief concerning Hendrix's date of birth to the further belief that he was a Virgo. As Lone$^+$ turns her attention to other matters, the two occurrent beliefs become standing beliefs—with the relevant pieces of information being stored in the neurologically integrated hard drive. Later Lone$^+$ realizes that Jimi Hendrix *died* on September 18 and was instead born on November 27. When Lone$^+$ acquires and endorses this new information, the information that Jimi Hendrix was born on September 18 is automatically overwritten on the hard drive and replaced by the information that he was born on November 27. Likewise, the information that he was a Virgo is overwritten and replaced by the information that he was a Sagittarius—assuming, as we may, that Lone$^+$ believes that Sagittarius are born between November 23 and December 21. In this way, Lone$^+$'s device is appropriately sensitive to evidence.

Second, Lone$^+$'s neurologically integrated search and storage device seems to integrate well with desires and other beliefs when it comes to the explanation of action. Since information-retrieval from Lone$^+$'s device hard drive is phenomenologically indistinguishable from information retrieval from biological memory, belief-desire explanations of action will have the right kind of structure. Indeed, when it comes to belief-desire explanations of action in cases like Lone$^+$, we need not make reference to any complex second-order beliefs. E.g., if Lone$^+$ were to search for the 1979 winner of the Danish 1st Division in soccer, the search would not need to be mediated by any second-order beliefs about the hard drive and the information on the hard drive. Rather, Lone$^+$ would ponder "Who won the Danish 1st Division in the 1978-1979 season?", and the answer would come to her in the same way as it would, had it been retrieved from biological memory. In this way, Lone$^+$ is like Inga when she recalls that the MoMA is on 53rd Street and completely *unlike* Otto, as described by Wikforss. Both Lone$^+$ and Inga have a standing belief—sitting in respectively a neurologically integrated hard drive and biological memory—waiting to become occurrent.[26]

The neurological device described in Lone$^+$ thus satisfies C_1-C_4 as well as Wikforss' additional condition C_W. Even by Wikforss more demanding standards, this device can hence sustain beliefs. To be sure, the device is not yet actual. So if we grant Wikforss her condition C_W, we may also grant her the conclusion that there is *currently* no device or resource that supports extended beliefs. That is, we may grant her that (CE1) and (CE2) type cognitive extensions do not explode—simply because there is no extension of belief and knowledge of these types in the first place. But, as highlighted, this makes Wikforss' reply to cognitive bloat non-principled and contingent on the current state of our technological development. If sufficiently advanced neuromedia become available—and Musk, Lynch, and Prichard all seem to think that they will—then Wikforss has not managed to rule out the feasibility of various forms of restricted omniscience through cognitive extension.

---

[26] The answer to Lone+'s question is, of course, "Esbjerg Forenede Boldklubber"—or "EFB" for short. During the 1978-1979 season EFB accumulated 46 points, claiming the championship with a convincing 6-point margin and a goal difference of 29 goals (59 for, 30 against). One author strongly approves of the content of this footnote—the other less so.

5. Against cognitive bloat (II): degree-theoretic, variantist cluster-model functionalism

While Wikforss aims to avoid cognitive bloat by appeal to a set of necessary and sufficient conditions on cognitive integrataion, we may worry that any such strategy will follow the sad pattern of the post-Gettier literature: a seemingly endless series of proposals for new conditions on cognitive integration—intended to rule out cognitive bloat—each undermined by counterexamples involving increasingly more complicated and fanciful external or integrated technological devices. In this section, we will look at J. Adam Carter and Jesper Kallestrup's alternative approach to cognitive integration:

> Instead we contend that the best way to settle the question of whether an external resource is integrated into a cognitive system is to determine whether it plays a characteristic functional role within that system, where that role is to be spelled out in some way other than in terms of necessary and sufficient conditions. (Carter and Kallestrup 2019: 9)

On Carter and Kallestrup's approach, to determine whether an external resource is integrated into a cognitive system, we can determine whether an external resource, relative to a particular situation of cognitive integration, satisfies most of the platitudes that we associate with integration. So there is no *fixed* set necessary and sufficient conditions that will determine cognitive integration across the board. This builds into the Carter-Kallestrup approach a certain flexibility that they can put to use in differet cases where cognitive bloat is unwanted.

More precisely, the Carter and Kallestrup's view involves four core theses:

1. Cluster-model functionalism.
2. Degree-theoretic approach to cognitive integration.
3. Variantism about cognitive integration.
4. Distinction between metaphysical cognitive integration and epistemic cognitive integration.

Let us go through each in turn.

Cluster-model functionalism. There is an overarching cluster of (regimented and Ramsified) platitudes C_1, …, C_m associated with cognitive integration. However, these conditions are not individually necessary and jointly sufficient. Rather, for specific instances of cognitive integration, what matters is not whether an external resource satisfies all of C_1, …, C_m, but merely a *majority* of these conditions. Supposing that there are only three conditions, Carter and Kallestrup explain their cluster-model functionalism as follows:

> We can then say that for a cluster of integration conditions C_1, C_2 and C_3 (i.e. a disjunction of conjunctions of most of C_1, C_2, C_3), resource *r* is integrated into cognitive system *S* if and only if *r* satisfies C_1 and C_2, or C_2 and C_3, or C_1 and C_3. So, it could be that external resource r^* is integrated into *S* for Otto in virtue of satisfying C_1 and C_2, while biological resource *r* is integrated into *S* for non-Alzheimers patients in virtue of satisfying C_1 and C_3. The key point is that the *same cluster* of integration conditions applies to extended and non-extended cognition. Otto has no less a claim to be a case of memory than ordinary cases of biological memory. (Carter and Kallestrup 2019: 12)

Carter and Kallestrup want their form of functionalism to be a flexible form of functionalism. Following David Lewis, they thus take the functional role to be tied to a *disjunction* of conjunctions of *most* of the conditions. This implies what we might call *realizer flexibility* because

there are several ways for an external device, resources, or onboard capacities to fulfill the functional role. Thus Otto's notebook might satisfy the C_1 & C_2 disjunct of the disjunctively characterized functional role while Inga's biological memory might satisfy the C_1 & C_3 disjunct.

Degree-theoretic approach to cognitive integration. Carter and Kallestrup suggest that we regard cognitive integration as a graded rather than an all-or-nothing matter. Let C_1, …, C_m be the relevant integration conditions. Then an external resource is integrated into cognitive system *S* to degree *n* if the resource satisfies exactly *n* conditions among C_1, …, C_m (where $0 \leq n \leq m$). An external resource is *maximally* integrated into *S* if the resource satisfies all of C_1 to C_m (i.e. if $n = m$), and it is *minimally* integrated into *S* if it satisfies none of C_1, …, C_m (i.e. if $n = 0$). This degree-theoretic approach to cognitive integration likewise gives a certain kind of flexibility. It allows for finer-grained distinctions and comparisons of cognitive integration.[27]

Variantism about cognitive integration. Carter and Kallestrup think that there are at least three sources of variation for cognitive integration. First, the cognitive profile of the subject is a determiner of which conditions are relevant to cognitive integration. This is clear from Carter and Kallestrup's idea above that certain conditions may be relevant for someone who relies on a notebook while different conditions may be relevant for someone who relies on biological memory.

Second, whether we are dealing with belief-*preservation* or belief-*generation* may impact which conditions are relevant to cognitive integration. For instance, since Otto's notebook is a candidate for preserving non-occurrent beliefs, the condition of past endorsement is highly relevant to his situation. Otto forms occurrent beliefs when endorsing information that he comes across, and in virtue of entering the information into the notebook as a result of this endorsement, the notebook becomes a candidate for preserving the relevant belief when it is no longer occurrent for Otto. In cases of belief-generation, however, the condition of past endorsement may not apply.[28]

Third, different cases may call for different assignments of relative weight to the conditions that are relevant to cognitive integration. In some cases, all relevant conditions are assigned equal weight while in others, some conditions are assigned greater weight than others. For example, if we consider cases of extended beliefs based on memory, we might want to assign the past endorsement condition greater weight than the conditions of typical invocation, automatic endorsement, and availability as and when needed.[29]

Distinction between metaphysical cognitive integration and epistemic cognitive integration. Farter and Kallestrup distinguish between two kinds of cognitive integration. Metaphysical cognitive integration pertains to whether a given device, resource, or capacity is sufficiently integrated with a subject's cognitive process to generate or preserve *belief*. Epistemic cognitive integration pertains to whether a given device, resource, or capacity is sufficiently integrated with a subject's cognitive process to generate or preserve *knowledge*. The relationship between metaphysical cognitive integration and epistemic cognitive integration is one of asymmetric entailment. The set of conditions associated with the former is a proper subset of the set of conditions associated with the latter. Knowledge asymmetrically entails belief.[30]

To illustrate the distinction between metaphysical and epistemic cognitive integration, Carter and Kallestrup consider Keith Lehrer's well-known case of Mr. Truetemp:

> Suppose a person, whom we shall name Mr. Truetemp, undergoes brain surgery by an experimental surgeon who invents a small device which is both a very accurate thermometer and a computational device capable of generating thoughts. The device, call it a tempucomp, is implanted in Truetemp's head so that the very tip of the device, no larger than the head of

[27] Carter and Kallestrup (2020: 13).
[28] Carter and Kallestrup (2020: 15).
[29] Carter and Kallestrup (2020: 13, 15).
[30] Carter and Kallestrup (2020: 5-8, 16).

> a pin, sits unnoticed on his scalp and acts as a sensor to transmit information about the temperature to the computational system of his brain. This device, in turn, sends a message to his brain causing him to think of the temperature recorded by the external sensor. Assume that the tempucomp is very reliable, and so his thoughts are correct temperature thoughts. All told, this is a reliable belief-forming process. Now imagine, finally, that he has no idea that the tempucomp has been inserted in his brain, is only slightly puzzled about why he thinks so obsessively about the temperature, but never checks a thermometer to determine whether these thoughts about the temperature are correct. He accepts them unreflectively, another effect of the tempucomp. (Lehrer 1990: 162-163)

The device implanted in Truetemp's brain, Tempucomp, is a *neuromedium* in the sense defined earlier. Most epistemologists think that the temperature readings from the thermometer generate belief but deny that Truetemps's beliefs about temperature qualify as knowledge. Carter and Kallestrup agree with these verdicts and take the distinction between metaphysical and epistemic cognitive integration to help explain this. How so?

Tempucomp is sufficiently well integrated into Truetemp's cognitive process to generate belief. Since we are dealing with a belief-*generating* process, Carter and Kallestrup do not take the past endorsement condition C_4 to be relevant. Conditions C_1-C_3, however, are relevant and also satisfied in the case. Tempucomp is typically invoked, easily available as and when needed, and Truetemp automatically endorses its temperature readings. This explains why tempucomp generates beliefs.

Since tempucomp is reliable, Truetemp has true beliefs formed through a reliable process. Yet, he fails to possess knowledge. According to Carter and Kallestrup, the reason is that epistemic cognitive integration, which is sufficient for knowledge, is more demanding than metaphysical cognitive integration, which is sufficient for belief. They consider two additional conditions on epistemic cognitive integration:[31]

(C_5) *Existence endorsement:* the existence of the belief-producing or belief-preserving device in relation to one's cognitive architecture is endorsed.

(C_6) *Reliability endorsement:* the reliability of the belief-producing or belief-preserving device is endorsed.

By adopting C_5 we impose the demand that the relevant cognizer is aware of the source that produces or preserves belief. By adopting C_6 we go one step further and demand that the cognizer likewise endorses the reliability of the source. While neither condition seems reasonable as a requirement on devices that generate or preserve belief, they might be more reaonsable for devices that generate or preserve *knowledge*.

Truetemp fails to know the temperature because the device endorsement condition C_5 fails. That is, Truetemp does not endorse the existence of tempucomp quite simply because the existence of the device is unknown to him. The failure of C_5 implies the failure of C_6. Since Truetemp does not endorse the existence of the belief-generating tempucomp, he does not endorse *its* reliability either. Whichever of C_5 or C_6 is relevant to knowledge-sufficient epistemic integration, Truetemp's true, reliably produced beliefs hence fail to qualify as knowledge.

But why should we accept C_5 or C_6 as conditions on knowledge? Carter and Kallestrup propose C_5 as a serious candidate condition because it prevents the subject from exemplifying a certain kind of meta-ignorance or meta-incoherence about the origin of beliefs. Truetemp exemplifies meta-ignorance about the origin of his temperature-related beliefs: if he were to reflect on where his temperature-related beliefs come from, he would draw a blank. Suppose we modify the Truetemp case as follows: everytime a temperature-related belief is generated

[31] Carter and Kallestrup (2020: 15-16).

by tempucomp, the device likewise generates a false belief in Truetemp about why he has that specific temperature-related belief—for instance, the device might generate the false belief that Truetemp has the temperature-related belief because he just read an *external* (i.e. normal) thermometer. With this modification to the case, Truetemp would exemplify meta-incoherence, as he would have a mistaken higher-order belief about the origin of his temperature-related beliefs.[32]

In relation to C_6, Carter and Kallestrup suggest that considerations concerning knowledge and cognitive achievement lend support to C_6. According to what we might call the *achievement thesis*, knowledge involves a cognitive achievement that possesses greater value than any cognitive achievement that falls short of knowledge—including, for instance, justified true belief. Relatedly, a knower is worthy of greater credit than someone who has a merely justified true belief. Carter and Kallestrup's idea seems be that endorsement of the reliability of the belief-generating or belief-preserving source contributes towards making the holding of that belief a cognitive achievement greater than any epistemic standing that falls short of knowledge.[33]

With an understanding of the four core theses that constitute Carter and Kallestrup's view, let us now turn to the challenge of cognitive bloat. As seen above, each core thesis yields increased flexibility along some dimension, and Carter and Kallestrup utilizes this flexibility in their response to cognitive bloat. To see how, let us first introduce three cases that Carter and Kallestrup discuss:

> TELO: Telo has a normally functioning biological brain. Like Otto, but also like other individuals with properly functioning biological memory, he relies on information in the environment to help structure his life, and this includes the information in his Verizon phone book, which sits in his desk drawer by his telephone. When Telo needs to call someone, he looks up the number. For Telo, the phone book prevents him from having to memorise everyone's phone number, just in order to make a call. (Carter and Kallestrup 2019: 3)

> TELO*: Telo* is like Telo, with the exception that he carries the phonebook everywhere he goes.[34]

> MEMORAID: Otto* is fitted with a remarkable chip called the *Memoraid*, developed to circumvent the need for Alzheimer's patients to use manual notebooks for cognitive offloading. Once the *Memoraid* chip is installed in a subject's prefrontal cortex, it *seems* to one that one is storing one's information normally, in one's biological memory; though, information is actually stored on the *Memoraid* chip. Likewise, retrieving information from the *Memoraid* chip is subjectively indistinguishable from retrieving information from biological memory. Otto*'s family felt it would be best not to tell Otto* that the *Memoraid* chip was installed, and so the entire procedure was done while Otto* was sleeping. Otto* is thus completely unaware that he is relying on the *Memoraid*—he has never even heard of this technology—and continues to believe that his memorial process involves him storing information in, and retrieving information from his biological brain. (Carter and Kallestrup 2019: 7)

---

[32] Carter and Kallestrup (2020: 6, 15). See also their comments in note 18. Carter and Kallestrup's MEMORAID case (2020: 7) involves meta-incoherence about the origin of the subject's beliefs. We return to the MEMORAID case below.

[33] Carter and Kallestrup (2020: 16). In the extensive body of work on the value problem, it is widely assumed that knowledge is more valuable than any epistemic standing that falls short of knowledge. See, e.g., Carter and Jarvis (2012), Greco (2009), Kvanvig (2003), Pritchard (2007), and Pritchard et al. (2018). Greco is a prominent advocate of the idea that knowledge is a cognitive achievement (2003, 2010, 2013). Another prominent advocate is Sosa (2007). Both Greco and Sosa likewise endorse the *ability intuition*, i.e. the idea that *S*'s true belief that *p* can only qualify as knowledge if the belief is due (at least to a considerable extent) to *S*'s cognitive ability. Furthermore, they take achievement and ability to be intimately related: knowledge is a cognitive achievement because it is due to ability.

[34] Carter and Kallestrup (2020: 14).

Telo is an example of *cognitive offloading* where a subject relies on an external device or resource in order to reduce the workload of an onboard capacity—in this case biological memory. The Telo* case shares fundamental features with our Lone case but lacks the feature of past endorsement. In both cases, the external resource is typically invoked, easily and readily available as and when needed, and information retrieved from the resource is automatically endorsed. However, while Lone's document contains only information that she has cut and pasted into the document, all entries in Telo*'s phone book were already there when he acquired it. MEMORAID is a neuromedia case that brings together elements from the Otto and Truetemp cases. Like Otto, Otto* is an Alzheimer's patient who needs to rely on external resources or devices in order to structure and navigate his daily life. Like Truetemp, Otto* has, unbeknownst to him, a neuromedium implanted. The neuromedium helps him structure and navigate his daily life.

What do Carter and Kallestrup say about the extent of beliefs and knowledge in Telo, Telo*, and Otto*?

The idea that Telo knows each entry in the phone book seems clearly wrong. The information is *just sitting there*, and Telo has never bothered to check most of the phone book entries. If knowledge were granted in the case of Telo, the same verdict would have to be given in relevantly similar cases. That is, ordinary subjects would know swaths of truths included in encyclopedia, newspapers, data bases, and many other kinds of resources—whether online or offline *and* whether they are even readily available as and when needed. So on pain of an extreme, and unwanted explosion of knowledge, we want to deny that Telo knows each entry in the phone book.

Carter and Kallestrup's view can yield this verdict. Telo does not know each entry in the phone book because the phone book does not satisfy the relevant conditions for *metaphysical* cognitive integration. As demanded by C_3, the phone book is not easily and readily available as and when needed. After all, it is sitting in the drawer of Telo's desk and, so, is not accessible when Telo is out. Hence the phone book does not even satisfy the conditions for belief and hence it cannot qualify as knowledge either.[35]

What about Telo*? Telo* carries the phone book with him wherever he goes, so it *is* readily and easily available as and when needed. The Telo* case, unlike the Telo case, thus satisfies C_3. Yet, the idea that Telo* knows each entry in the phone book still militates against our ordinary thinking about knowledge. If knowledge were granted in the case of Telo*, it would have to be granted in all relevantly similar cases. So any ordinary subject who is constantly connected to the internet through her smartphone would qualify as knowing swaths of unaccessed truths on Wikipedia, the news, journals, and many other kinds of resources. Thus, again, on pain of an unwanted extreme explosion of knowledge, we want to deny that Telo* knows each entry in the phone book.

Carter and Kallestrup aim to secure this result by claiming two things. First, the type of knowledge relevant to Telo* is the external device or resource counterpart of memory-based knowledge. The external resource stores information and thereby relieves Telo* of having to memorize phone numbers. In this way, the resource is meant to serve as a proxy for memory. Second, since the relevant resource is meant to be belief-*preserving*, the condition C_4 of past endorsement is crucial to yield a degree of cognitive integration that is sufficient for belief. However, Telo* has never endorsed all entries in the phone book, and so the phone book cannot be said to preserve a belief for each of its entries. Consequently, there is no explosion of knowledge in Telo* because there is no explosion of *belief*.[36]

---

[35] Carter and Kallestrup (2020: 14). They also argue that Telo does not satisfy the condition of typical invocation (i.e. C_1). We leave this aside for present purposes.

[36] Carter and Kallestrup (2020: 14-15).

Finally, what about Otto*? Carter and Kallestrup grant that the neurologically integrated *Memoraid* chip satisfies the conditions C_1-C_4 for metaphysical cognitive integration. The chip, recall, plays the exact same role as biological memory, and by design, storing and retrieving information from it is phenomenologically indistinguishable from storing and retrieving information from biological memory. The chip is typically invoked, and stored information is readily and easily available as and when needed. Furthermore, the information is automatically endorsed upon retrieval and is there because Otto* endorsed it in the past.[37]

While Carter and Kallestrup grant Otto* beliefs of the information on the *Memoraid* chip, they deny that these Memoraid-based beliefs qualify as knowledge. Reason: they take Otto*'s situation to be, in all epistemically relevant respects, like Truetemp's situation. So even though Truetemp and Otto*'s beliefs are true and preserved or generated by a reliable device, both lack relevant knowledge. The theoretical backing for this diagnosis is again found in the distinction between metaphysical and epistemic cognitive integration. The specific conditions for epistemic integration require endorsement of the existence of the belief-preserving or belief-generating source (C_5), or the reliability of this source (C_6). Otto* fails to satisfy both of these conditions because, like Truetemp, Otto* is unaware that a device has been implanted. Thus he cannot endorse the existence of Memoraid, let alone its reliability. Consequently, the conditions for epistemic integration are not met—although those for metaphysical integration are—in which case Otto*'s Memoraid-based beliefs do not qualify as knowledge.

6. Against: against cognitive bloat (II)

We do not think that Carter and Kallestrup's view can deal with the issue of cognitive bloat or knowledge explosion in a principled manner—and very much for the same reason as we gave in the discussion of Wikforss' proposal. Even if certain conditions *currently* block extended belief or knowledge in a given range of cases, there is no guarantee that they will continue to do so as available technologies become more advanced and sophisticated. That is, although Carter and Kallesteup give examples of how (CE1), (CE2), and (CE3) types of cognitive extension need not cause cognitive bloat, the specific bloat-preventing conditions (C_5) and (C_6) are, as we shall see, vunerable to more sophisticated (CE3)-type technological extensions.

We will proceed along two paths. We first assess the effectiveness of Carter and Kallestrup's response given the conditions on metaphysical and epistemic cognitive integration that they explicitly state and consider. Given these conditins, we argue that their view fails to deal with cognitive bloat or restricted omniscience in the key cases we have presented. We then consider the possibility of adding further conditions on knowledge to block this argument. We argue that there is considerable further work to be done before this strategy can be executed. Committing to this further work, we suggest additionally, may also come with a substantial philosophical commitment to a view that we label "cluster-epistemology".

Recall that Carter and Kallestrup's conditions (C_5) and (C_6) provide theoretical backing for the intuitions that Truetemp and Otto* lack relevant knowledge. Appealing to conditions (C_5) and (C_6) can help block cognitive bloat in cases where

1. unbeknownst to some subject *S*, a device is neurologically implanted in *S*;
2. the device feeds *S* mostly true information;
3. when fed information by the device, the subject forms a corresponding belief; *but*
4. *S* has no inkling about the source of the belief (or, alternatively, incorrectly thinks that it came about in a certain way).

[37] Carter and Kallestrup (2019: 7-8).

To illustrate, let us consider a case that has the features described in 1-4. The case brings together elements from the cases of Truetemp and our Lone$^+$ case:

> *Jennifer*
> Unbeknowst to Jennifer, a neurological device is implanted in her brain. The device feeds her the correct answer to any query or question she might have (whenever a correct answer is possible). The phenomenology is such that Jennifer feels like the answer just 'pops up' in her mind. (Alternatively, we can suppose that the phenomenology is indistinguishable from recall from biological memory—and so, that Jennifer mistakenly thinks that the answer was retrieved from memory.) Jennifer spends a weekend thinking about different topics—among others, all birthdates of past Academy Award winners, British prime ministers, and French presidents; the first 10,000 entries in the multiplication table of 17; and the daily conversion rate from US dollars to Danish kroner for the past 20 years. Given the nature of the implanted device, by the end of the weekend, Jennifer has formed a wide range of true beliefs through a reliable belief-generating process.

Since Jennifer endorses neither the existence of the neurological device nor its reliability, conditions C_5 and C_6 are not satisfied in the case, and hence Jennifer's device-generated beliefs all fail to qualify as knowledge.

This is all well and good: C_5 and C_6 can do *some* work. Either condition on its own can deal with Truetempesque cases of restricted omniscience. Yet, the force of C_5 and C_6 is limited. In the key cases we have considered, the subject's epistemic situation is *unlike* that of Truetemp, Jennifer, and Otto*.

Consider again Lone and Lone$^+$. Lone herself has created the document on her smartphone and she only includes information in the document that comes from sources that are reliable and which she takes to be reliable. She is thus very much aware of the existence of the information-sustaining source and the sources from which the stored information has been taken. Moreover, Lone recognizes that her way of collecting, storing, and retrieving information is indeed reliable. As for Lone$^+$, she was knowingly fitted with a neurological storage device, and she carried out the internet searches that led to information being collected and stored. Lone$^+$, too, is thus aware of the existence of the information-sustaining or preserving source and the sources from which the stored information has been drawn. Like Lone, Lone$^+$ also recognizes that her way of collecting, storing, and retrieving information is reliable.

These reminders about our key cases bring out something interesting: in addition to C_1-C_4, the cases of Lone and Lone$^+$ satisfy C_5 as well as C_6. Interestingly, then, the cases of Lone and Lone$^+$ satisfy all conditions explicitly stated and discussed by Carter and Kallestrup. Whether metaphysical or epistemic cognitive integration is in question, our two key cases thus exhibit the *highest degree* of cognitive integration: they satisfy *all* of C_1-C_6. Lone and Lone$^+$ must both be attributed a very wide range of beliefs, as their respective devices are metaphysically cognitively integrated in a way sufficient for belief. Among these many beliefs, furthemore, all true, reliably preserved ones wil qualify as knowledge, as their respective devices are epistemically cognitively integrated in a way sufficient for knowledge. So we can conclude that the Lone case *still* shows that ordinary subjects are in a position easily to attain various forms of restricted omniscience. Nothing that Carter and Kallestrup have said or argued undermines the potency of the case—as well as the more futuristic Lone$^+$ case.

Of course, most ordinary subjects do not spend one month cutting and pasting like Lone does. But if we grant the extended mind thesis, it seems that ordinary subjects actually *do* possess a considerable body of extended knowledge. This is because many ordinary people collect, endorse, and store information on their smartphones—in some cases using services such as

Dropbox—and because smartphones easily satisfy C_1-C_4. If we worry about this conclusion, the question is of course whether we can use Carter and Kallestrup's conditions C_5 and C_6 to block it. Arguably, we cannot. For people endorse the existence of smartphones (sic!), and even do so specifically as devices on which information can be stored and from which information can be retrieved. So C_5 is satisfied for these mundane cases. Condition C_6 is satisfied too. While some people endorse explicitly the reliability of the process of storing and retrieving information from their phones, others do it implicitly. At least that would best explain why so many people continuously rely on their phones for several mundane everyday purposes.

So even taking on board C_5 and C_6 as conditions on extended knowledge, the upshot is that many forms of restricted, but surprising omniscience are within the epistemic reach of ordinary subjects. If you worry about cognitive bloat, you will thus presumably continue to worry after having familiarized yourself with Carter and Kallestrup's propsal.

Let us next offer a few words on neuromedia in Carter and Kallestrup's framework. While Carter and Kallestrup reject the idea that Truetemp and Otto* (in MEMORAID) possess knowledge, their rejection does not seem to be rooted in a principled opposition to neuromedia's being sufficiently integrated. If Lone+'s neurological devices are sufficiently advanced for knowledge-sufficient epistemic cognitive integration, might Carter and Kallestrup then concede that Lone+ possesses knowledge? The following passage suggests that they do:

> Of course, we can additionally stipulate that TrueTemp also satisfies (5) and (6) [i.e. C_5 and C_6] such that he would have such knowledge in virtue of possessing an integrated, cognitive ability involving such device. In that case, however, we lose grip on the intuition that TrueTemp mysteriously acquires extraordinary knowledge. Instead the case illustrates an enhancement of epistemic powers which our common-sense functionalist would welcome. (Carter and Kallestrup 2019: 17)

If Carter and Kallestrup grant that Lone+ enjoys various forms of restricted omniscience, they in effect concede that cognitive bloat or knowledge explosion gets realized in certain circumstances.

Adding conditions C_5 and C_6 to C_1-C_4 cannot address the issue of cognitive blot in a principled way. For what it is worth, as just seen, it may be that Carter and Kallestrup do not wish to do so. Although they do not address cases like Lone where no fancy neuroimplants are invoked, this is arguably the most promising line for them to take. First, it aligns most naturally with the permissive nature of the functionalism at the heart of their position. Anything that plays the right kind of functional role can be integrated sufficiently into a cognitive process to generate or sustain belief or knowledge. Second, if Carter and Kallestrup want to argue that Lone or Lone+ fails to possess vast amounts of extended knowledge, they must put one or more additional conditions on the table that can serve to underwrite a no-knowledge verdict in the relevant case(s). Executing this task leaves considerable work to be done, and it may bring to light an easily overlooked but very substantial philosophical commitment. Let us see why.

Since Carter and Kallestrup seem to be somewhat sympathetic to the idea that Lone+ is restrictedly omniscient, let us focus on Lone. The task of supporting a no-knowledge verdict in the case of Lone requires executing (at least) the following two sub-tasks:

1. Specify a condition on epistemic cognitive integration that Lone's document fails to satisfy.

2. Provide a motivation for adopting the condition specified as a constraint on epistemic cognitive integration.

While the requirements of the first sub-task are obiovus, the requirements of the second sub-task ensure that ad hocery is avoided. Unless it enjoys some degree of independent motivation, it would be ad hoc to introduce a condition not satisfied by Lone's document for the sole purpose of ruling it out as a knowledge-preserving resource.

Carter and Kallestrup offered worries about meta-incoherence or meta-ignorance as motivations for adopting C_5. Since subjects like Truetemp are completely clueless about the origins of the relevant beliefs, they are meta-ignorant. Other subjects might take their beliefs to have a source that, in fact, they do not have. Such subjects are meta-incoherent. C_5 rules out unfortunate meta-predicaments of this sort. However, since Lone is very much aware of the nature of the relevant belief-preserving process, it does not seem plausible that a "Lone-excluding" new condition can be motivated by appeal to meta-incoherence or meta-ignorance.

C_6 was pitched as a way of putting flesh to the idea that knowledge is a cognitive achievement over and above any epistemic standing that falls short of knowledge. Carter and Kallestrup understood this idea in explicitly axiological terms, meaning that knowledge has greater value than any epistemic standing that falls short of knowledge.[38] If the new Lone-excluding condition is to be motivated by considerations of achievement—as linked to value—some story must be told about how and why satisfaction of the condition contributes value. This would be a significant item on the to-do list.

The so-called *ability* and *anti-luck intuitions* are other potential motivations for candidate conditions on knowledge-sufficient epistemic integration. According to the ability intuition, knowledge requires that the subject's beliefs, at least to a considerable extent, are due to her cognitive abilities. Some virtue epistemologists—including Greco and Sosa—take ability and achievement to be intimately related: knowledge is a cognitive achievement *because* it is due to ability.[39] According to the anti-luck intuition, knowledge is incompatible with epistemic luck. Safety and sensitivity are widely regarded as different ways to give body to the anti-luck intuition. Put informally, safety is the idea that the subject's belief could not easily be wrong. With slightly more rigour: if *S* were to believe that *p*, then *p* would be true. Put informally, sensitivity is the idea that the subject's belief needs to be sensitive or attuned to the facts. With slightly more rigour: if *p* were false, then *S* would not believe that *p*.[40]

While there are, of course, other motivations for conditions on knowledge on the market, Carter and Kallestrup are somewhat non-committal when it comes to specific epistemological theories. Their chief concern seems to be to present their framework—i.e. degree-theoretic, variantist cluster-model functionalism—at a somewhat structural or abstract level and carry out their discussion of cognitive integration at this level. Thus located, their framework is compatible with conditions that track different kinds of epistemological motivations. The only definite commitment to be gleaned from Carter and Kallestrup's discussion is negative: they are *not* de facto reliabilists. This is clear from their treatment of the Truetemp case. As for positive commitments, they seem sympathetic—but not definitely committed—to the achievement intuition.

The search for an additional condition om epistemic cognitive integration raises a substantial issue. Once a degree-theoretic, variantist cluster-model functionalist has committed to specific conditions on epistemic integration, the question will arise how similar or dissimilar the epistemological motivations are for the various conditions in the framwork. This will raise the further issue whether some incarnations of cluster-model functionalism carry a commitment to *cluster-model epistemology*. For it may well turn out that the motivations for the various conditions incorporated into some incarnation of the cluster-model functionalist framework get

[38] Carter and Kallestrup (2020: 16).
[39] Greco (2003, 2010, 2013), Sosa (2007).
[40] Prominent advocates of sensitivity include Dretske (1971) and Nozick (1981). Prominent advocates of safety include Pritchard (2005), Sosa (1999), and Williamson (2001).

associated with different epistemological theories. In that case, cluster-model functionalism will be accompanied by cluster-model epistemology: a framework that builds in conditions on epistemic integration that track ideas typically thought to be distinctive of different—and competing—epistemological theories.

We articulate and highlight the potential commitment to cluster-epistemology here for two reasons. First, the potential commitment is easily overlooked—especially if, as in Carter and Kallestrup's own work, the discussion of cluster-model functionalism is carried out at a somewhat abstract, structural level. Second, cluster-epistemology would make for a very significant commitment. Why? Because it amounts to a substantial philosophical view that many epistemologisits likely will find controversial, and some even unpalatable.

In sum, if Carter and Kallestrup want to give a no-knowledge verdict in the case of Lone, significant work remains to be done. They must specify a condition that can do the job and offer a motivation for this condition. Once specific conditions are added to the cluster-model functionalist framework, the issue of a potential commitment to cluster-model epistemology must be considered, together with the potential ramifications of such a commitment. Again, it may be that Carter and Kallestrup accept that Lone's document is sufficiently epistemically integrated to qualify as a knowledge-preserving resource. However, let us remind ourselves that this would amount to a concession of one of our main points: namely that ordinary subjects are in a position easily to attain various forms of restricted omniscience. If you worry about cognitive bloat, you will thus continue to worry.

7. Concluding remarks

Above we have offered a detailed discussion of the cognitive bloat issues surrounding the extended mind thesis and its epistemic counterpart. We argued that neither Wikforss nor Carter and Kallestrup can satisfactorily avoid cognitive bloat. In light of our discussion, people who worry about cognitive bloat might be disappointed by the prospects of extended epistemology. And to a certain extent, we agree. Neither Wikforss nor Carter and Kallestrup seem to be opposed to the functionalist core that underlies the extended mind thesis. And this, we believe, might be the central reason why neither of their proposals can sustain a principled rejection of cognitive bloat—in particular as cognitive bloat is manifested by restricted omniscience.

It is part of the basic nature of functionalist views—indeed, one of its underlying motivations—not to be particularly discriminatory. After all, the key idea behind functionalism about a given phenomenon *X* is that anything that fulfills the functional role of *X* gets to *be X*. Thus, if you grant functionalism about belief, then anything that plays the functional role of belief gets to be belief. This is so whether information sits in biological memory, in non-onboard resources and devices, or neurologically integrated in the brain. As we saw in our discussion of the Lone$^+$ case, if one beefs up the functional role of belief by imposing an additional constraint regarding sensitivity to evidence and direct interaction (i.e. Wikforss' condition C_W), information in seamlessly integrated computer devices can still play the functional role of belief. Similarly, although Carter and Kallestrup endorse a refined form of functionalism, we argue above that information in both external computer devices (Lone) and seamlessly integrated computer devices (Lone$^+$) satisfy the conditions for metaphysical and epistemic integration. Thus, given functionalism, the proposals of Wikforss and Carter and Kallestrup cannot completely eliminate the risk of cognitive bloat.

In tracing what seems to be problematic about cognitive bloat, considerations of counterintuitiveness play a central role—as we saw in section 1. Interestingly, neuromedia seem highly relevant to the issue of counterintuitiveness. For is it really counterinutive that we

can easily acquire and store a wealth of beliefs and knowledge through such devices? Here it should be borne in mind that the envisioned neuromedia allow for *seamless* integration in that storing and retrieving information from their hard drive is phenomenologically indistinguishable from storing and retrieving information from biological memory. After Lone$^+$ has spent a month looking up, endorsing, and storing all kinds of information, the phenomenology of retrieving that information will be the same as for biological memory.

We are inclined to think that the air of cognitive bloat associated with neuromedia can have various sources—reflecting, ultimately, theoretical commitments that may be articulable once the issues of extended belief, knowledge, and cognitive bloat have been investigated further. For someone like Wikforss, the counterintuitiveness of smartphone-generated cognitive bloat is grounded in the reaction that the information stored on smartphones is not sufficiently integrated to qualify as belief and hence as knowledge. This reaction was theoretically unpacked as C_W—the evidence sensitivity and direct interaction condition—and imposed as an additional constraint on belief. However, this suggests that once neuromedia allow for seamless integration, the source of counterintuitiveness vanishes.

But, even given seamless integration, some might claim that it is counterintuitive that we can easily come to believe and know many things through neuromedia. In this case, the counterintuiveness may ultimately be grounded in a fundamental opposition to the very idea that *any* non-onboard resource or device can sustain belief or knowledge. Such opposition would go naturally with a rejection of functionalism. If functionalism, ultimately, is the source of the counterintuitiveness of cognitive bloat, the underlying view will hold that non-onboard resources and devices—whether computerized or not, and whether seamlessly integrated or not—can have neither qualitative nor quantitative significance in relation to belief and knowledge. Quite simply put: there is no such thing as extended belief and, thus, no such thing as extended knowledge either.

So perhaps, at the end of the day, a discussion of cognitive bloat centers around a discussion of functionalism, and, as such, around a discussion of the very intelligibility of the extended mind thesis. We will have save this discussion for another day.

References

Allen-Hermanson, S. 2013. Superdupersizing the Mind: Extended Cognition and the Persistence of Cognitive Bloat. *Philosophical Studies* 164: 791–806.

Bjerring, J. C. and N. J. L. L. Pedersen. 2014. All the (Many, Many) Things We Know: Extended Knowledge. *Philosophical Issues* 24: 23–38.

Block, N. and J. Fodor 1972. What Psychological States Are Not. *Philosophical Review* 81: 159–181.

Carter, J. A. and B. Jarvis. 2012. Against Swamping. *Analysis* 72: 690–699.

Carter, J. A., B. Jarvis and K. Rubin. 2015. Varieties of cognitive achievement. *Philosophical Studies* 172: 1603–1623.

Carter, J. A. and J. Kallestrup. 2020. Varieties of cognitive integration. *Noûs* 54: 867–890.

Carter, J. A. and D. Pritchard. 2019. *Journal of Philosophy and Medicine* 44: 220–242.

Chalmers, D. J. 2008. Foreword. In Clark 2008, ix–xvi.

Clark, A. 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.
———. 2010. Memento's Revenge. In *The Extended Mind*, ed. R. Menary, 43–66. Cambridge, MA: MIT Press.

Clark, A. and D. Chalmers. 1998. The Extended Mind. *Analysis* 58: 7–19.

Dretske, F. 1971. Conclusive Reasons. *Australasian Journal of Philosophy* 49: 1–22.

Farkas, K. 2012. Two Versions of the Extended Mind Thesis. *Philosophia* 40: 435–447.

Goldman, A. 1979. What Is Justified belief? In *Justification and Knowledge*, ed. G. Pappas, 1–23. Boston: D. Reidel.
———. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

Greco, J. 2003. Knowledge as Credit for True Belief. In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, eds. M. DePaul and L. Zagzebski, 111–134. Oxford: Oxford University Press.
———. 2009. The Value Problem. In *Epistemic Value,* eds. A. Haddock, A. Millar and D. Pritchard, 313–321. Oxford: Oxford University Press.
———. 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.
———. 2013. A (different) virtue epistemology. *Philosophy and Phenomenological Research* 85: 1–26.

Kelp, C. 2013. Extended cognition and robust virtue epistemology. *Erkenntnis* 78: 245–252.

Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.

Ludwig, D. 2015. Extended cognition and the explosion of knowledge. *Philosophical Psychology* 28: 355-368.

Lynch, M. P. 2014. Neuromedia, extended knowledge and understanding. *Philosophical Issues* 24, 299–313.
———. 2016. *The Internet of Us: Knowing More and Understanding Less in the Age of Big Data*. New York: W. W. Norton.

Musk, E. 2019. An Integrated Brain-Machine Interface Platform with Thousands of Channels. *Journal of Medical Internet Research* 21: e16194. DOI: 10.2196/16194

Nozick, R. 1981. *Philosophical Explanations*. Oxford: Oxford University Press.

Palermos, O. 2011. Belief-Forming Processes, Extended. *Review of Philosophy and Psychology* 2: 741–765.
———. 2014. Knowledge and cognitive integration. *Synthese* 191: 1931–1951.
———. 2015. Active externalism, virtue reliabilism and scientific knowledge. *Synthese* 192: 2955–2986.

Pritchard, D. 2005. *Epistemic Luck*. Oxford: Oxford University Press.
———. 2007. Recent Work on Epistemic Value. *American Philosophical Quarterly* 44: 85–110.
———. 2010. Cognitive Ability and the Extended Cognition Thesis. *Synthese* 175: 133–151.
———. 2018a. Neuromedia and the Epistemology of Education. *Metaphilosophy* 49: 328–349.
———. 2018b. Pritchard, D. 2018b. Extended Knowledge. In *Extended Epistemology*, eds. J. A. Carter, A. Clark, J. Kallestrup, O. Palermos and D. Pritchard, 90–104. Oxford: Oxford University Press.
———. 2018c. Extended virtue epistemology. *Inquiry* 61: 632–647.

Pritchard, D., J. Turri and J. A. Carter. 2018. The Value of Knowledge. In *The Stanford Encyclopedia of Philosophy*, ed. E. Zalta (Summer 2018 Edition). URL = <https://plato.stanford.edu/archives/sum2018/entries/knowledge-value/>.

Putnam, H. 1967. Psychological Predicates. In *Art, Mind, and Religion*, eds. W. H. Capitan and D. D. Merrill, 37–48. Pittsburg: Pittsburgh University Press.

Rubert, R. D. 2004. Challenges to the Hypothesis of Extended Cognition. *Journal of Philosophy* 101: 389–428.

Shah, N. and D. Velleman. 2005. Doxastic deliberation. *Philosophical Review* 114: 497–534.

Smart, P. 2012. The web-extended mind. *Metaphilosophy* 43: 446–463.
———. 2018. Emerging Digitial Technologies: Implications for Extended Conceptions of Cognition and Knowledge. In *Extended Epistemology*, eds. J. A. Carter, A. Clark, J. Kallestrup, O. Palermos and D. Pritchard, 266–304. Oxford: Oxford University Press.

Sosa, E. 1999. How to Defeat Opposition to Moore. *Philosophical Perspectives* 13: 141–154.
———. 2007. *A Virtue Epistemology*. Oxford: Oxford University Press.

Sprevak, M. 2009. Extended Cognition and Functionalism. *Journal of Philosophy* 106: 503–527.
———. 2019. Extended Cognition. Routledge Encyclopedia of Philosophy Online. DOI: 10.4324/9780415249126-V049-1

Wikforss, Å. 2014. Extended Belief and Extended Knowledge. *Philosophical Issues* 24: 460–481.

Williamson, T. 2001. *Knowledge and Its Limits*. Oxford: Oxford University Press.