



# On some intracranialist dogmas in epistemology

J. Adam Carter<sup>1</sup>

Received: 25 October 2021 / Accepted: 25 August 2022  
© The Author(s) 2022

## Abstract

Research questions in mainstream epistemology often take for granted a *cognitive internalist* picture of the mind. Perhaps this is unsurprising given the seemingly safe presumptions that (i) knowledge entails belief (*viz.* the *entailment thesis*) and that (ii) the kind of belief that knowledge entails supervenes exclusively on brainbound cognition. It will be argued here that (*contra* orthodoxy) the most plausible version of the entailment thesis holds just that knowledge entails *dispositional* belief. However, regardless of whether occurrent belief supervenes only as the cognitive internalist permits, we should reject the idea that dispositional belief supervenes only in cognitive internalist-friendly ways. These observations, taken together, reveal two things: first, that a cognitive internalist picture of the mind is much more dispensable in epistemology than has been assumed and, second, that pursuing questions in *extended epistemology* need not involve any radical departure from the commitments of more traditional epistemological projects.

**Keywords** Extended epistemology · Knowledge and belief · Cognitive internalism, Dispositional belief · Extended cognition

## 1 Introduction

It is rare that we find epistemologists (as opposed to philosophers of mind<sup>1</sup>) defending specific views about the metaphysical nature of knowledge-apt representational states. However, mainstream work in epistemology—especially on the nature of knowledge—almost invariably<sup>2</sup> *presupposes* a certain background picture of cognition that fits very well with the idea that knowledge is necessarily stored and generated in the head.<sup>3</sup>

<sup>1</sup> For discussion, see, e.g. Cummins (1996), Lycan (2000), Ramsey (2017) and Shea (2013).

<sup>2</sup> For some exceptions, see, e.g. Palermos (2018), Carter (2013), Palermos and Pritchard (2013) and Bjerring and Pedersen (2014).

<sup>3</sup> The kind of presupposition here is pragmatic presupposition (in the sense of, e.g. Stalnaker 1973)—in that both sides of the disputes about the nature of knowledge act as though it is in the common ground between them that, e.g. beliefs are in the head and memory processes are in the head.

✉ J. Adam Carter  
adam.carter@glasgow.ac.uk

<sup>1</sup> University of Glasgow, Glasgow, Scotland

In the philosophy of mind and cognitive science, this background picture has a name—*cognitive internalism*—the view that, necessarily, cognition supervenes on brainbound, biological properties of the cogniser.<sup>4</sup> Cognitive internalism is, as Fred Adams and Kenneth Aizawa (2009) put it, tantamount to a dictum of common-sense—viz. that the ‘mind is in the head’.

When articulated as a view about cognitive *processes* (i.e. memory storage and retrieval), cognitive internalism is usually read as maintaining that, necessarily, cognitive processes (e.g. memory storage and retrieval) play out entirely inside the head<sup>5</sup>; alternatively, such processes are materially realised exclusively by physical processes in the brain. When framed as a view about *states* of cognition (e.g. beliefs), the view implies that your beliefs are literally in your head, in the sense that the physical subvenient bases of your beliefs are all and only intracranial subvenient bases.<sup>6</sup>

Here is perhaps the most straightforward picture of how cognitive internalism is so easily ‘smuggled in’—and uncritically so—as a presupposition in epistemology. Propositional *knowledge*—of central interest in epistemology—is assumed *ex ante* to *entail* belief, truth and justification, as per the traditional JTB analysis.<sup>7</sup> The project of analysing knowledge—which dominated the second half of twentieth century epistemology—aimed to work out how these three (and perhaps other) conditions relate to each other when one has knowledge. *Belief* relates to the other two conditions—at least, in a way that matters for analysing knowledge—in so far as beliefs are propositional attitudes with a representational (i.e. mind to world) direction of fit. That is, after all, what really matters in the analysis of knowledge, because it is exactly *this kind of a thing* that is capable of being true and justified and thus, as the thought goes, capable of being known. And furthermore, in at least *paradigmatic* cases of propositional knowledge (think of simple perceptual knowledge—viz. your knowledge that there is a hand in front of you, which you generate and then retain in memory), it seems plain enough that *brainpower* is going to be both necessary and sufficient to (i) generate the (occurrent) propositional attitude with a representational direction of fit (i.e. *that* there is a hand in front of you right now) and then to (ii) store it, as a dispositional belief, in memory.

Thus, the ‘map’ to unearthing the cognitive internalist presupposition in mainstream thinking about knowledge is accordingly a pretty direct one, with two key ‘links’ in the chain: the first that gets us from knowledge to belief (via the ‘entailment thesis’ that knowledge entails belief) and the second that gets us from belief to cognitive internalism (where the latter is the assumed picture about how the former is realised and maintained).

<sup>4</sup> For some prominent recent defences of cognitive internalism against challenges from embedded and extended cognition camps, see Adams and Aizawa (2008; 2010).

<sup>5</sup> See, e.g. Carter et al. (2014), Wheeler (2018), Carter et al. (2016), Kiverstein (2018), Palermos (2018), Palermos and Pritchard (2013) and Pritchard (2010).

<sup>6</sup> The thesis applies not only to beliefs, but also to, e.g. desires and emotions. Thus, a view on which emotions supervene partly on something extracranial, including on partially extracranially driven appraisal processes (see, e.g. Carter et al., 2016; Kruger and Szanto 2016), is incompatible with cognitive internalism.

<sup>7</sup> For some detailed overviews of this project, see Shope (2017) and Ichikawa and Steup (2018). For criticism, see Williamson (2002, Ch. 1).

In what follows, I am going to challenge both links of this chain. Or, more carefully, I am going to show that despite initial appearances, there is really no plausible construal of the first link (that knowledge entails belief) that should pressure epistemologists who care about knowledge to assume a cognitive internalist picture of the mind.

Here is the plan. §2 clarifies the sense in which knowledge entails belief, and in doing so defends a dispositional reading of the entailment thesis. §3 then puts pressure on cognitive internalism generally, and then more specifically on the idea that such a view would explain any better than cognitive externalism the datum that knowledge entails dispositional belief. Taken together, §§2–3 challenge some of the dogmas of traditional epistemology that seem to block progress on certain questions in *extended epistemology*—e.g. epistemological questions whose answers make reference to transcranially supervenient cognitive processes and states.

## 2 Knowledge and belief

Whilst the idea that knowledge entails belief is widely assumed,<sup>8</sup> it is rarely argued for positively (apart from being defended against objections<sup>9</sup>), with two notable exceptions being G. E. Moore (1962) and Keith Lehrer (1968). Moore famously tried to show that knowledge entails belief via a (albeit somewhat odd) linguistic test, and Lehrer (1968) opted for a proof aimed at showing that knowledge formally entails belief. Neither is promising.

According to Moore:

There certainly is a common use of belief in which ‘I believe’ entails ‘I don’t know for certain’. Is there another in which ‘I know for certain’ entails ‘I believe’? One reason why it seems so is because ‘I thought I knew’ entails ‘I believed’ (1962, 115).

It does seem plausible that a speaker who says ‘I thought I knew’ that *p* is committed in some way to accepting that they believed that *p*. But let us simply *grant* for the sake of argument that patterns like the one Moore mentions constitute linguistic evidence that knowledge entails belief (either occurrent or dispositional). Even on this charitable assumption, there is, as Carolyn Black (1971) has observed, also linguistic data that would seem to support the very *opposite* conclusion. Take for example, this case: ‘I say that my books are in my office.’ You ask ‘Do you believe that your books are in your office?’ I say ‘*No*. I *know* that my books are in my office’ (Black, 1971, 155–56, my italics). The felicitousness of *this* kind of exchange is a problem for arguments that attempt to establish that knowledge

<sup>8</sup> In particular, we find this assumption in the decades of critical response to Gettier (1963). See, e.g. Shope (2017).

<sup>9</sup> Defences of the knowledge-belief entailment against objections have largely focused on responses in the late 1960s and 1970s to Radford’s (1966) ‘unconfident examinee’ case, which we discuss later in this section.

entails belief (of any sort) simply on the basis of our patterns of using the words ‘knows’ and ‘believes’.<sup>10</sup>

So what about Lehrer’s (1968) proof? Here is the proof, which he takes to be sufficient to establish to a doubter that knowledge entails belief.

1. If  $S$  does not believe that  $P$ , then  $S$  does not believe that he knows that  $P$ .
2. If  $S$  does not believe that he knows that  $P$ , then, even though  $S$  correctly says that  $P$  and knows that he has said that  $P$ ,  $S$  does not know that he correctly says that  $P$ .
3. If, even though  $S$  correctly says that  $P$  and knows that he has said that  $P$ ,  $S$  does not know that he correctly says that  $P$ , then  $S$  does not know that  $P$ .
4. (Therefore) If  $S$  does not believe that  $P$ , then  $S$  does not know that  $P$  (1968, 498).

There are problems with both premises (1) and (3). The problem with (1) is that it is either false or *at best* questionbegging, given what Lehrer was attempting to do here. Just consider that the kind of opponent Lehrer is out to convince might very well think that “ $S$  knows that  $p$ ” is compatible with the antecedent of (1). But then, (1) comes out false if  $S$  does not believe that  $p$  because  $S$  knows that  $p$ , given that, on that supposition, it is possible that  $S$  will believe that  $S$  knows that  $p$ . But even if this problem with (1) could be dealt with, there are independent problems with (3): just suppose your friend tells you they do not know all the lines of a certain poem by William Blake, but then (after telling you this) they proceed to recite the poem perfectly; this seems like a plausible case where—even though they do not know that they have correctly recited it—they nonetheless know the lines.<sup>11</sup> They had them mastered better than they had thought. This assessment of the case, however, is incompatible with (3).

<sup>10</sup> It is worth clarifying that the challenge to Moore here that appeals to Black’s linguistic data does not depend on Black’s linguistic data *lacking* any alternative felicitous reading. For instance, it might be pointed out that (in response to Black’s data) it could also be felicitous to challenge the assertion ‘No, I know that my books are in my office’ by pointing out that, perhaps, what the asserter meant with the denial ‘No’ is that they did not *merely believe* the books were in the office, not that they did not believe it (at all) whilst also knowing it. For my purposes here, I’m happy to grant that such a challenge would be a natural one, but to reiterate that the problem for Moore remains for the following reason: in so far as Moore is defending the entailment thesis on the grounds of their being felicitous use patterns of ‘knows’ and ‘believes’ that would be explained by the thesis, data such as Black’s registers that there are *also* felicitous patterns of use that would count against the entailment thesis. In sum, then, relying simply on felicitousness to support the entailment thesis is not going to be decisive. Thanks to a reviewer for suggesting further clarification on this dialectical point.

<sup>11</sup> For a similar case, see Black (1971, 157). Granted, some cases with a similar structure might be seem more plausibly interpreted as cases where a defeater is present which undermines the first-order knowledge. However, in cases like this—where it looks like knowledge that is not only possessed but manifest (i.e. in the recitation of the lines)—a ‘knowledge defeat’ assessment is not clearly more plausible than a reading on which one is simply mistaken *whether* one knows. Further cases featuring practical knowledge can help illustrate the point further. Consider, for example, a variation on a case due to Setiya (2008), where one does not know whether the anaesthesia from a hand surgery will wear off by noon and so does not know whether they can clench their fist at noon when they try. Even so, as the thought goes, one upon clenching one’s fist at noon (when the anaesthesia in fact has worn off, even if luckily so), one can know that they are clenching their fist when doing so intentionally and successfully at noon. In such a case, as with our Blake case, it looks like one’s manifesting one’s knowledge in action is compatible with their despite having been unsure whether they have the relevant knowledge. Such cases would seem to put pressure on (3) of Lehrer’s argument *even if* we grant that some cases of failing to know whether one knows that  $p$  constitute a defeater for one’s knowledge that  $p$ . Thanks for a referee for suggesting further discussion on this point.

Interestingly, we do not find many other attempts<sup>12</sup> to positively establish the widely held assumption that knowledge entails belief (and, as it turns out, Lehrer himself abandoned his own proof later,<sup>13</sup> opting instead for a view on which knowledge entails not belief but *acceptance*).

Instead, what much of the literature on the knowledge-belief ‘entailment thesis’ concerns is whether outlying attempts to *challenge* the thesis are sound. The most widely discussed case on this score—also one that involves (contentiously<sup>14</sup>) a kind of ‘knowledge-with-lack-of-confidence’ structure—is due to Colin Radford (1966):

UNCONFIDENT EXAMINEE: Kate is taking a history test. She had studied carefully and has been doing well on all the questions so far. She has now reached the final question, which reads “What year did Queen Elizabeth die?” As Kate reads this question she feels relief, since she had expected this question and memorized the answer. But before Kate can pause to recall the date, the teacher interrupts and announces that there is only one minute left. Now Kate panics. Her grip tightens around her pen. Her mind goes blank, and nothing comes to her. She feels that she can only guess. So, feeling shaken and dejected, she writes “1603”—which is of course exactly the right answer.<sup>15</sup>

As David Rose and Jonathan Schaffer (2013) put it, ‘The case of *Unconfident examinee* represents the leading challenge to the orthodox idea that knowledge entails belief’ (2013, 20). Apart from this classic case from the mid 1960s—and the extensive critical response to it (on both sides), which fizzled out in the 1980s—the most notable recent lines of argument against the idea that knowledge entails belief, all in the past 10 years, are due to Blake Myers-Schulz and Eric Schwitzgebel (2013), Katalin Farkas (2015) and Susanna Schellenberg (2017a, 2017b). Myers-Schulz and Schwitzgebel present experimental evidence<sup>16</sup> that the intuition in UNCONFIDENT EXAMINEE that Kate has knowledge without belief (that Queen Elizabeth died in 1603) is robust, and on this basis, purport to give a ‘second wind’ to the old counterexample to the orthodox presumption that knowledge entails

<sup>12</sup> Whilst Armstrong’s (1969) paper *Does Knowledge Entail Belief* is ostensibly a defence of the claim, it is less an attempt to establish the thesis than it is to defend it against cases like those from Radford (1966).

<sup>13</sup> See Lehrer ([1990] 2018).

<sup>14</sup> As I suggest below, I am dubious that the case shows what Radford initially took it to show; critical assessment of the case will be useful nonetheless, as we will see closer. Consideration of the case indicates sense in which knowledge (contra Radford) *does* involve belief, albeit, of a dispositional variety.

<sup>15</sup> This is Rose and Schaffer’s succinct characterisation of what is a longer dialogue used to capture the example by Radford (Radford 1966, 2–3).

<sup>16</sup> For some additional experimental evidence that is meant to vindicate the idea that UNCONFIDENT EXAMINEE counts against the entailment thesis, see Murray et al. (2013).

belief.<sup>17</sup> Farkas, on the other hand, uses cases of *extended cognition* (e.g. cases where one offloads one's memory tasks to a notebook or a smartphone) as plausible cases where one has knowledge without belief, albeit, knowledge stored externally. Finally, Schellenberg's tack is to cast doubt on whether the entailment thesis holds specifically in cases of perceptual knowledge, where (arguably) one knows simply via *seeing* that something is so, and regardless of whether one forms a belief.

Just as we have seen that Moore and Lehrer did not plausibly *demonstrate* that knowledge entails belief (either occurrent or dispositional), there is also a good case to be made that *none* of the above attempts aimed at establishing that knowledge *does not* entail belief succeed, at least in so far as none of these strategies plausibly demonstrate that knowledge does *not* entail dispositional belief. This point turns out to be relevant to the wider transition from 'knowledge to belief, and then from belief to cognitive internalism', given that occurrent belief rather than dispositional belief is more *prima facie* plausibly wed to a cognitive internalist picture of the mind.

Regarding the UNCONFIDENT EXAMINEE case, the pressure against the entailment thesis is really the strongest when we contrast (i) the observation that Kate's lack of confidence in the proposition that the Queen died in 1603 does not seem to preclude her from knowing it, and indeed, manifesting that knowledge unconfidently, with (ii) the thought that Kate must believe and thus *consciously endorse* the content that the Queen died in 1603 at some time if she is to know it at that time. The force of UNCONFIDENT EXAMINEE against the entailment thesis lies in the fact that it leads us to embrace (i), and then on that basis reject (ii).

But importantly, a rejection of (ii) is *compatible* with the thesis that knowledge entails belief, so long as 'belief' is understood in a dispositional sense, where dispositional beliefs are merely *available to mind for endorsement*<sup>18</sup> even when the content of a dispositional belief is not (occurrently) consciously endorsed.<sup>19</sup> Rose and Schaffer (2013) support this rationale on the basis of two considerations. First, Kate's memory trace<sup>20</sup> (viz. that Queen Elizabeth died in 1603) is not destroyed. Second, her guess is no accident.<sup>21</sup> On the second point, they write:

Indeed it seems as if her memory trace must still be not just present but actually operating in the background to guide her actions, even if she is unable in the moment to appreciate the fact. Putting these two reasons together—to the

<sup>17</sup> As they note: "A majority of respondents ascribed knowledge [...] whilst only a minority ascribed belief" (Myers-Schulz and Schwitzgebel 2013, sec. 3).

<sup>18</sup> For some explicit discussions of dispositional belief and its relationship with occurrent belief, see Armstrong (1973), Lycan (1986) and Audi (1994). For an overview, see Schwitzgebel (2019, sec. 2.1).

<sup>19</sup> Though, perhaps as Murray et al. (2013) maintain, one must have assented to the proposition at some point in the past.

<sup>20</sup> Memory traces (sometimes referred to as 'engrams' in psychology) are taken to be the means by which we store memories in the brain. For a recent overview of work on memory traces, see de Brigand (2014). For philosophical discussion of memory traces in the epistemology of memory, see Bernecker (2010).

<sup>21</sup> The idea that one might know via reliable guessing, even when one lacks confidence, is given an explicit defence in Ernest Sosa's (2015) virtue epistemology. In particular, see Sosa's case of the eye examination (in his 2015, Ch. 3).

extent that it is useful to operate with the picture of a “belief box” in which various propositions are stored—we find it natural to think of Kate as having the proposition that Queen Elizabeth died in 1603 lodged in her belief box throughout. She stored it there during her studies and is still unconsciously guided by it when she “guesses.” Indeed we find it natural to imagine that—perhaps later that very day—Kate will recover from her panic and recall the information readily enough. She has the information stored in mind. She is merely temporarily blocked from accessing it normally (2013, 24–25).

It looks, then, as though we should deny that Kate has a dispositional belief only if we are prepared to say that her temporary block is permanent rather than temporary. But *even if it were* permanent, note that the kind of block she has just prevents her from accessing the information stored in mind *normally*. It does not prevent her from accessing it *at all* for the reason that this information stored in memory continues to guide her actions. Of course, *were it* to somehow be blocked off from even doing *that*, then we might then deny her the dispositional belief, on account that it is inaccessibly stored in memory. However, on that kind of a scenario, there would then be no pressure to attribute to her knowledge. If she guessed correctly, it would be by sheer luck.

The above considerations cast Myers-Schulz and Schwitzgebel’s (2013) experimental results in a different light. From the fact that folks are more likely to attribute knowledge than belief in UNCONFIDENT EXAMINEE, we have no good reason to reject the entailment thesis—at least not without a clearer sense of which sense of the polysemous ‘belief’ the participants took themselves to be withholding whilst at the same time attributing knowledge. Interestingly, as more recent experimental studies indicate,<sup>22</sup> when the same experiments are run whilst eliciting the dispositional reading of belief more so than it was elicited in the original experiments, people’s intuitions no longer disproportionately line up with attributing knowledge rather than belief.

The take-away lesson from the UNCONFIDENT EXAMINEE case seems to be this: the case (i) purports to show that it is not the case that knowledge entails belief; (ii) it plausibly *does* demonstrate that knowledge does not entail occurrent belief; but (iii) it *does not* succeed in showing that knowledge *does not* entail dispositional belief—on the contrary, we would plausibly be less likely to attribute knowledge in the case were dispositional belief *not* present.

Although Farkas’s argument against the knowledge-belief entailment thesis is ostensibly very different from the line of argument that proceeds from the THE UNCONFIDENT EXAMINEE case, an appreciation of Farkas’s wider argument shows that it ultimately slots into the very same kind of general (i, ii, iii) structure.

Her argument takes as its basis a case of cognitive offloading from memory to notebook. The case—involving the characters ‘Otto’ and ‘Inga’—was originally used by Clark and Chalmers (1998) as an argument *against* cognitive internalism,

<sup>22</sup> These are the results reported by Rose and Schaffer (2013), who replicated the Myers-Schulz and Schwitzgebel experiments whilst more explicitly eliciting the dispositional reading of ‘belief’.



and in favour of the idea that cognition can extend beyond the boundaries of the skull and skin. The Otto and Inga case—and the thesis of ‘extended cognition’ more generally—will be discussed in some detail in §3. For our purposes now, though, let us focus squarely on how Farkas thinks the case supports a rejection of the orthodox idea that knowledge entails belief.

The key first step for Farkas is to take a queue from Edward Craig’s (1991) thinking about the *purpose* of the concept of knowledge, an understanding of which Craig thinks would help to illuminate what falls in its extension.<sup>23</sup> According to Craig: ‘[k]nowledge is not a given phenomenon, but something that we delineate by operating with a concept which we create in answer to certain needs, or in pursuit of certain ideals’ (1991, 2). On Craig’s view, we ‘create’ the concept of knowledge in order to meet the need we have to *flag reliable informants*. And so, on the Craigian view, the *function* of the concept of knowledge is to flag reliable informants, and relatedly, an appreciation of this function as the function it is should guide our thinking about what falls within the extension of the concept of ‘knowledge’.

Now, with these Craigian ideas assumed, Farkas encourages us to think about the case of Otto and Inga:

OTTO AND INGA: Inga would like to go to the Museum of Modern Art (MoMA); she recalls that the MoMA is on 53rd street, and she sets off accordingly. Otto suffers from severe memory loss and therefore he keeps all important information recorded in a notebook which he carries with him all the time. When he decides to go to MoMA, he looks up the whereabouts of the museum, finds it’s on 53rd street, and then he sets off. Many people agree that Inga had had the belief that the Museum of Modern Art was on 53rd street even before the issue came up in connection to her current visit. But Clark and Chalmers claim that if Inga has the belief, so does Otto, even before he looked up the information in his notebook. Otto has reliable, constant and easy access to the contents of his notebook, and he endorses the contents of his notebook automatically. This, according to Clark and Chalmers, is enough to qualify him as having the belief (2015, 190).

Farkas’s own idiosyncratic take on this case fits neither with traditional thinking (on which Otto neither believes *nor* knows that MoMA is on 53rd street in virtue of storing this information as he does in his notebook but not in his head) but *nor* does Farkas’s assessment line up with the point Clark and Chalmers originally used the case to make, which is that—as they see it—Otto’s memory (and thus, his dispositional beliefs stored in memory) lies partly in the notebook, external to his biological brain. Farkas thinks—and we need not get in to the details just yet, but we’ll return to them—that we should agree with the traditionalist that the cognitive differences between Otto and Inga are substantial enough that, when it comes

<sup>23</sup> This ‘Craigian’ idea that the nature of knowledge is something we can fruitfully illuminate by first inquiring into what the concept of knowledge is for—viz. what the function of the concept of knowledge is—has enjoyed some more recent support under the heading of ‘function-first’ epistemology. See, e.g. McKenna (2013) and Hannon (2018). Cf., Gerken (2015).



to attributing ‘belief’, we should do so *disanalogously*, to Inga but *not* to Otto. On the other hand, however, she thinks we should part ways with the traditionalist—and simply be guided by Craig—when it comes to assessing whether to attribute *knowledge* to Otto. Recall again the Craigian idea that the point of the concept of knowledge is to track reliable informants, and just consider in this light how we use ‘knowledge’ to track such informants in cases of, e.g. seeking phone numbers. As Farkas writes, in ‘some everyday contexts, it is very natural to attribute knowledge to subjects who are in Otto-type situations. You ask me if I know NN’s phone number, and I say “sure”, reaching for my smartphone’ (Farkas, 2015, 190).

Putting this all together, Farkas thinks we have compelling reason to think Otto *knows but does not believe* that MoMA is on 53rd street, and, *a fortiori*, that the knowledge-belief entailment thesis is false. Now, I have suggested at the outset that I think Farkas’s argument ends up slotting into the (i, ii, iii) structure that characterised the purported argument against the entailment thesis from UNCONFIDENT EXAMINEE. I now want to explain why.

First, consider that one tempting spot to challenge Farkas’s reasoning is her claim that Otto and Inga are different enough that we should *not* attribute dispositional belief across the cases symmetrically. Why not? Why *are not* Clark and Chalmers right about this, as opposed to the traditionalist? Fortunately, there is a way press back against Farkas without fully opening that can of worms (we will circle back to it in §3), which is to suggest that *by her own lights* we ought to attribute Otto a dispositional belief. The reasoning here is that attributing such a dispositional belief is the most promising way for Farkas to vindicate her claim that Otto has (extended) *knowledge*. Farkas’s rationale for attributing Otto extended knowledge, after all, is meant to be guided by the Craigian idea that we should use ‘knowledge’ to track reliable informants. The presumption here (which we may grant, *ex hypothesi*) is that Otto is such a reliable informant; ask him where MoMA is, he can reliably tell you (via a process that involves consulting his notebook rather than biomemory). Now, *what is it that grounds Otto’s reliability about where MoMA is?* It is hardly a brute fact that he is reliable. On the contrary, he is a reliable informant because he *reliably stores the information* (just like Inga does); his information is correct, easily available for endorsement, etc. Indeed, it thus looks quite a bit like the thesis that Otto knows where MoMA is (in virtue of what is in his notebook) would be explained—even granting the Craigian story—by his having something that looks an awful lot like a dispositional belief.

Now, a traditionalist has at this juncture might try to dig their heels in as a matter of principle: ‘Cognitive internalism is true and so, necessarily, all cognition plays out in the head; therefore, Otto simply *can’t* have a dispositional belief externally stored’. But—crucially—it looks like this kind of a principled reason is already out the window for Farkas, who explicitly allows *knowledge* outside the head. Farkas’s line that Otto’s case features knowledge without belief accordingly occupies a curious area of dialectical space: her claim that Otto has knowledge (that MoMA is on 53rd street) itself seems best *explained* by his having a dispositional belief, in virtue of how he stores the information he does, not in biomemory, but in the notebook. Farkas, of course, denies that he has a dispositional belief (by appealing to cognitive internalist thinking), *but* that denial would itself be principled denial only if Farkas

were to *also* deny that he has extended knowledge (which she of course does not deny).

Putting this all together, then, it looks as though—as with UNCONFIDENT EXAMINEEE—the case of Otto and Inga (at least, as Farkas is using it) exhibits (i, ii, iii) structure; it is a case that Farkas (i) purports to use to show that it is not the case that knowledge entails belief; (ii) the case plausibly *does* demonstrate that knowledge does not entail occurrent belief (given that Otto clearly lacks such a belief, no less than Kate does in UNCONFIDENT EXAMINEE), but (iii) it *does not* succeed in showing that knowledge *does not* entail dispositional belief, and, if anything, only serves to positively reinforce this idea.

Let us round out our discussion of the knowledge-belief entailment thesis with a brief look at Susanna Schellenberg's domain-specific dismissal of the idea that knowledge entails belief. The line she advances is 'domain specific' because it is meant to apply exclusively to perception, and thus to perceptual knowledge. According to Schellenberg's view of perceptual knowledge, *capacitivism*, a subject (*S*), has perceptual knowledge that *p* by *seeing* that *p*, which requires that *S* employ 'a capacity to single out what she purports to single out' (2017, 318) and *S*'s mental state (whereby *S* sees that *p*) must have 'the content it has in virtue of *S* having successfully employed her capacity to single out what she purports to single out' (2017, 318).

As the reader will have noticed, 'belief' does not feature in the above story. This, Schellenberg thinks, is just as it should be. She writes:

Orthodoxy has it that one cannot know that *p* without believing that *p*. Capacitivism is neutral on whether there is any such belief condition on knowledge. This is attractive, since arguably, we know that *p* simply in virtue of seeing that *p*. By contrast, we do not believe that *p* simply in virtue of seeing that *p*. After all, I can see that *p* without forming any beliefs (2017, 318).

Of course, even if Schellenberg is right, she will have been right about a story of perceptual knowledge acquisition. What about perceptual knowledge *retention*? Suppose you see that *p* at  $t_1$ . At  $t_2$ , you are no longer thinking about *p*. But, if someone asks you at  $t_2$  what *p* looked like, you remember and can tell them. But this would turn out to be mysterious if at  $t_2$  you did not retain this information about *p* in a way that was then later available to mind for endorsement.<sup>24</sup> But that is just the mark of a dispositional belief. To the extent that Schellenberg's capacitivism is a correct story of perceptual knowledge acquisition, this story looks to be compatible with the version of the entailment thesis that has seemed most plausible so far—viz. that knowledge requires at least dispositional belief.

Recall now that the 'map' to unearthing the cognitive internalist dogma in mainstream thinking about knowledge had two key 'links' in the chain, one from knowledge to belief (Link 1), the other from belief to cognitive internalism (Link 2).

<sup>24</sup> I am using the simplified idea of 'available to mind for endorsement' from Rose and Schaffer (2013, secs. 1.3, S22).

This section—critically examining Link 1—reveals that the most charitable way to unpack Link 1 is as:

- Link 1<sub>dispositional</sub>: propositional knowledge → (entails) dispositional belief rather than:
- Link 1<sub>occurrent</sub>: propositional knowledge → (entails) occurrent belief:

However, from Link 1<sub>dispositional</sub>, we most plausibly get to cognitive internalism only by way of:

- Link 2<sub>dispositional</sub>: dispositional belief → (is best explained by) a cognitive internalist picture of the mind rather than:
- Link 2<sub>occurrent</sub>: occurrent belief → (is best explained by) a cognitive internalist picture of the mind.

But this is where the overarching story—from mainstream thinking about knowledge to the cognitive internalist assumption that tacitly underlies it—begins to show some real cracks. Just consider that, whereas Link 2<sub>occurrent</sub> is *prima facie* very plausible (if not obvious to many), Link 2<sub>dispositional</sub> really is not.

The reason Link 2<sub>occurrent</sub> seems platitudinous is that occurrent belief is usually taken to involve *consciously* entertaining (and subsequently endorsing) a proposition, and a biological brain is plausibly (though this point is of course debatable) necessary *and* sufficient for this kind of conscious activity. Accordingly, it is *prima facie* plausible that cognition *of the sort that is realised exclusively as the cognitive internalist countenances* is what furnishes us with whatever occurrent beliefs we have. Crucially, however, a biological brain is—though obviously sufficient—*not necessary* for realising the kind of thing that hosting a dispositional belief is generally taken to involve, which is the *storing* of information that is *available* to us for conscious endorsement. If anything, the ubiquity of cognitive offloading suggests that even though biological brains suffice for storing information available for conscious endorsement, they are obviously not necessary because we use them *for this very purpose* increasingly less—especially when it comes to practical information of the sort we rely on to structure our lives. It is, then, at best *prima facie* plausible that cognition of the sort that is realised exclusively as the cognitive internalist countenances furnishes us with only some of our dispositional beliefs. But this means, then, that the phenomenon of dispositional beliefs is best explained by a picture of the mind that allows for storage of information available for conscious endorsement to sometimes be handled intracranially, *sometimes (and increasingly often) not*.

At this juncture, the proponent of cognitive internalism might simply double down as follows: ‘even if the sense in which knowledge entails belief is best understood as Link 1<sub>dispositional</sub> *rather* than Link 1<sub>occurrent</sub>, and indeed even if it *looks* as though we can make sense of many of the dispositional beliefs we have without assuming anything like cognitive internalism, it remains that cognitive internalism stands up as an independently and overwhelmingly plausible ‘pillar’ in the philosophy of mind. It establishes the bounds of cognition in a way that aligns with centuries of philosophical thinking, and we are better placed simply accepting the implications of cognitive internalism wherever they lead us,

*even where* they don't align so well with our other commitments (at least, when these other commitments lack the kind of 'bedrock' status that cognitive internalism enjoys). And so, despite initial appearances to the contrary, we should not accept but resist the temptation to think that the process of storing information available for endorsement in a notebook or iPhone (rather than in biomemory) is a genuine *cognitive* process, and thus, we should resist attributing 'beliefs' and 'knowledge' on the basis of such storage.'

Does the proponent of cognitive internalism here have a point? This really depends on whether cognitive internalism is (or deserves to be) the kind of 'pillar' in our theorising that the above reasoning suggests. As it turns out, pillars fall, and, lately, old 'internalist pillars' in particular have been falling right and left.

### 3 Objection and reply

Thus far, the line advanced is that a cognitive internalist picture of the mind is much more dispensable in epistemology than has been assumed. However, let us now take seriously the following kind of rejoinder: *regardless of whether* cognitive internalism is indispensable (or not) in our epistemological theories, it independently enjoys a kind of sacrosanct status as a theory of the metaphysics of mind; thus, as the thought would go, we should be taking it for granted *anyway* in epistemology, given the strong independent reasons for thinking that cognitive internalism is on entirely solid ground.

Let us think through this line of critique. Until relatively recently, the study of knowledge was—following a tradition inherited from Descartes<sup>25</sup>—a thoroughly 'internalist' enterprise in three key ways.

First, it used to be taken for granted that the content of our thoughts is determined entirely by the inner workings of the mind—viz. *content internalism*.<sup>26</sup> On this way of thinking, your intentional attitudes (e.g. your beliefs and other attitudes that are *about* things) are about the things that they are about (rather than about other things) in virtue of your psychological states and nothing else. Any two people in the same psychological states, then, must be thinking *about* the very same thing. For those (like Descartes) who are aligned with this kind of thinking, it is easy to see how 'rigorous philosophical inquiry must proceed via an inside-to-out strategy', and of course, as was apparent in the *Meditations*, from *this* kind of methodological starting point, the challenge of (non-circularly<sup>27</sup>) defeating the sceptic becomes especially difficult.<sup>28</sup>

<sup>25</sup> The typical reference point here is the *Meditations*; however, Descartes' internalist picture of the mind and the way it represents the world is not limited to his epistemology; it is also central to his wider philosophy of mind. See, e.g. Cottingham (2002).

<sup>26</sup> For some representative discussions of cognitive internalism, see Loar et al. (1988), Kriegel (2013) and Fodor (1987).

<sup>27</sup> As Descartes suggested, even from a content internalist starting point, one can 'transition' from knowledge of one's mind to knowledge of the world if one is entitled to the claim that there is a non-deceiving God. However, a famous objection to Descartes is that it is not clear how one can get to this conclusion non-circularly. For discussion, see, e.g. Markie (1992).

<sup>28</sup> Arguably, as some epistemic externalists (e.g. Sosa 1997) have pointed out, analogous problems arise for indirect realist strategies in the epistemology of perception (e.g. Moore's) which purport to vindicate external world perceptual knowledge as based on inference from information just about the qualitative character of our experiences.

Even so, content internalism is not itself an epistemological thesis (even if it has some epistemological ramifications); it is a thesis about how the content of our thoughts and words is individuated. An importantly different kind of internalism—also inherited from Descartes and widely assumed until around the 1970s<sup>29</sup>—is *epistemic internalism*.<sup>30</sup> Epistemic internalism is not a thesis about what our thoughts and words refer to, but about what kinds of things *justify* our beliefs in a way that matters for knowledge. It is in principle compatible with either content internalism or content externalism.<sup>31</sup> What the epistemic internalist maintains is that epistemic justification is *solely* determined by factors that are internal<sup>32</sup> to a person.<sup>33</sup> Such factors include, e.g. what mental states one is in and what is accessible to one via reflection alone. A simple reason why this kind of view (a centre-piece of Descartes' epistemology, but with origins as early as the *Theatetus*) has plausibly enjoyed the support it has is that we tend to think of the kind of justification that matters for knowledge as being associated with reasons and evidence, and the matter of what reasons and evidence one has seems—on the face of things—to be determined by factors internal to one (e.g., what your mental states are).<sup>34</sup>

Rounding out the three internalist 'pillars' of Cartesian epistemology is our old friend *cognitive* internalism on which what is claimed to be 'internal' to a thinker is not the content of their thoughts (content internalism) or what matters for justifying their beliefs (epistemic internalism), but rather the *material realisers* of their cognising, including whatever thoughts and beliefs they have, justified or not.

The suggestion—canvassed in the previous section on behalf of the traditionalist—that cognitive internalism is some kind of unalterable 'pillar' that must not be dislodged is really not very compelling in the context of appreciating that—of these three internalist 'pillars of Cartesianism'—the *first two have already fallen*, and both within just the past 50 years.

Content externalists in the 1970s<sup>35</sup> and 1980s<sup>36</sup> have shown how our environments play a crucial role in individuating meaning and mental content, and to such

---

<sup>29</sup> It is important not to run together the longstanding endorsement of an internalist picture of epistemic justification with the related, but separate, issue of whether this picture of epistemic justification has a longstanding place in a justified-true-belief analysis of propositional knowledge. Although Dutant (2015) has recently called received thinking about the place of the JTB analysis in the history of epistemology since Descartes into doubt, this doubt does not apply to the largely internalist way in which epistemologists have (until the rise of externalism in the 1960s and 70 s) thought and talked about knowledge-relevant justification.

<sup>30</sup> For a sample of epistemic internalist positions in epistemology, see Alston (1988), Chisholm (1973), Conee and Feldman (2004) and Huemer (2006).

<sup>31</sup> Though, for some critical discussion on this point, see Chase (2001), Pritchard and Kallestrup (2004) and Carter et al. (2014).

<sup>32</sup> The 'internal' in internalism is usually taken to be something like 'internal to one's psychology' or to one's 'mental states'. And *those* are things that are almost invariably understood as brainbound. That said, it is at least in principle possible to envision a more radical kind of epistemic internalism, paired with a more inclusive conception of what one's psychology and mental life can consist of. For a discussion of this more radical kind of spin on epistemic internalism, see Carter and Palermos (2015).

<sup>33</sup> See, e.g. Poston (2014) and Madison (2017).

<sup>34</sup> A good example of this kind of assimilation of 'reasons' and 'evidence' talk with epistemic internalism is found in Chisholm (1977).

<sup>35</sup> (Putnam 1975).

<sup>36</sup> (Burge 1986).

an extent that content internalists are nearly extinct in 2021. As Juhani Yli-Vakkuri and John Hawthorne (2018) put it, in a recent monograph purporting to be the final nail in the coffin of this kind of internalist thinking, ‘entanglement of our minds with the external world runs so deep that no internal component of mentality can easily be cordoned off’. With the exception of Hawthorne and Yli-Vakkuri’s purported final takedown, content externalism is now so popular it is rarely taken to need any additional argument. Essentially, philosophical thinking has ‘flipped’ almost completely since the mid 1970s, and on a position fundamental to our grip on the very nature of thought.

What about *epistemic internalism*, then? It has slowly but steadily (since the 1960s) been heading the way of content internalism. According to results from a PhilPapers Survey published by David Chalmers and David Bourget in 2014, only about a quarter of 931 philosophers surveyed (246/931 (26.4%)) self-identify as epistemic internalists.<sup>37</sup> This is so even though internalism captured the default position in epistemological theory from Plato, to both rationalists (Descartes) and empiricists (Locke and Hume)<sup>38</sup> all the way up to Gettier (1963). Whilst debates between epistemic internalists and externalists remain contentious, one thing that is clear is that epistemic internalism is no longer the default view but rather the exception.<sup>39</sup>

So is *cognitive internalism* then the only ‘Cartesian pillar’ that should be thought of as ‘safe’ from the externalist wave—and as such, permanently fixed? The short answer is ‘no’ for the reason that this final internalist pillar has at least partially (arguably: *mostly*) fallen *already*, as the past 20 years of the philosophy of cognitive science suggests. It is just that—put simply—this news has not quite spread to mainstream epistemology.

The cracks in cognitive internalism started quite small.<sup>40</sup> Forget iPhones and the like for a moment, and just think about your hands, and how you move them around, gesturing as you talk. This kind of gesturing not only facilitates communication, but also helps language *processing* (McNeill, 1992). Likewise, consider the baseball outfielder (e.g. McBeath et al., 1995) trying to catch a fly ball, by running in a direction that makes the ball appear to follow a straight line. In doing this, the outfielder is solving a complex problem not just by perception, but by a kind of ‘perception–action coupling’—viz. by using perceptual information to guide movement and then using movement to hold the perceptual information constant.

<sup>37</sup> See also <https://survey2020.philpeople.org/> for an updated version of the survey.

<sup>38</sup> For a clear presentation of Hume’s internalist foundationalism, see, e.g. Sosa (1980). For Locke (e.g. *Essay IV*, xvii, 24), epistemic internalism was a feature of his wider assimilation of epistemic justification with doing one’s epistemic duty.

<sup>39</sup> Key to epistemic internalism’s downfall is arguably the sheer strength of the thesis itself—viz. its contention that *everything* that matters for justification must be internal to one’s psychology, or available to one by reflection alone. As externalist epistemologists such as Alvin Goldman (1979) have emphasised, such views cannot countenance the insight that the reliability of a belief-forming process is amongst those things that seem to matter. But, the externalist is not similarly restricted; the externalist can consistently allow that some of what matters for epistemic justification is determined by factors internal to one’s psychology, though not everything.

<sup>40</sup> Outside of analytic philosophy, the idea that cognition might be embodied was already gaining some traction in the nineteenth century continental phenomenology of perception. For an overview of the history of embodied cognition, see Gallagher (2014).

The above are just some representative example cases—others (many of which have appeared just since the 1990s<sup>41</sup>) involve visual consciousness,<sup>42</sup> concepts,<sup>43</sup> memory,<sup>44</sup> moral cognition,<sup>45</sup> etc.—which have been taken to favour the view that cognition is best understood as not only taking place in the brain, but more widely, as *embodied*. Wilson and Foglia (2017) articulate the core of the ‘embodied cognition’ thesis as follows:

Many features of cognition are embodied in that they are deeply dependent upon characteristics of the physical body of an agent, such that the agent’s beyond-the-brain body plays a significant causal role, or a physically constitutive role, in that agent’s cognitive processing.

What the evidence for embodied cognition suggests is that cognition is not merely (as traditionalists would have it) ‘sandwiched between, while segregated from’ perception and action. The dependence of the former on the latter is simply too deep to separate in the clean way the traditionalist/internalist would want.<sup>46</sup>

Getting down to brass tax: if *any* part of an agent’s *non-brain* body has *ever* played a physically constitutive role in cognitive processing, then strictly speaking, cognitive internalism is false.<sup>47</sup> And as more evidence has come in that validates this very idea, embodied cognition has increasingly taken over as the ‘default’ position in cognitive science. As Fred Adams (2010)—a committed traditionalist—concedes: ‘The view that cognition is embodied [...] is rapidly gaining prominence in the world of cognitive science, *and is aiming for dominance* (2010, 619). According to Lawrence Shapiro (2014), embodied cognition is ‘now one of the foremost areas of study and research in philosophy of mind, philosophy of psychology and cognitive science’.

It is hard to see how cognitive internalism should deserve any kind of sacrosanct status when the tide in cognitive science is now generally against it. But if *that* is right, then is not it just a clear mistake for epistemologists to cling tacitly to cognitive internalism?

Maybe—in a sense—not. Consider this line of argument: ‘Let’s assume cognition is embodied – granted! Even so, this is a far cry from suggesting that you can have *beliefs* in your *phone*. Your phone is *not* part of your brain *or* your biological body!’ This kind of rejoinder suggests that perhaps the best candidate for a plausibly sacrosanct thesis in the neighbourhood of cognitive internalism is not strict cognitive

<sup>41</sup> For an overview, see Gibbs (2005) and Wilson and Foglia (2017, sec. 5).

<sup>42</sup> See, e.g. Noë (2005) and Hurley (1998).

<sup>43</sup> (Lakoff 2012).

<sup>44</sup> (Sutton 2006).

<sup>45</sup> (Haidt, Koller, and Dias 1993).

<sup>46</sup> This ‘inseparability’ idea is sharpened further by those friends of embodied cognition who go a step further to think of the body and mind as a kind of dynamical system (e.g., Chemero 2011; Beer 1995; S. Orestis Palermos 2016, 2014).

<sup>47</sup> See Gibbs (2005). Some cognitive scientists have called into question whether the empirical evidence supports what proponents of embodied cognition take it to support. For a response to some of these rejoinders, see Miracchi (2021).



internalism after all, but rather the more permissive cognitive *bio-internalism*—the thesis that cognition is essentially *biologically* realised.

However, even if we shift the goal posts of sanctity from cognitive internalism to *cognitive biointernalism*, we still fail to capture anything properly sacrosanct. Two straightforward challenges on this score come from cognitive neuroscience over the past 6 years alone: (i) the 2015 creation of the first artificial neuron (Simon et al., 2015) and (ii) the first successful case (in 2019) of creating artificial memories from scratch and implanting them in mice, where the artificial memories guided behaviour indistinguishably from non-implanted memories (Vetere et al., 2019). Note that in neither of these cases is cognition realised entirely as the biointernalist would have it.

A larger elephant in the room, however, comes from the philosophy of cognitive science, where researchers are increasingly open and explicit in their denial of even *cognitive biointernalism*. It is here where it will be useful to circle back to the case of Otto and Inga, originally due to Clark and Chalmers (1998). Both the cognitive internalist *and* the more permissive cognitive biointernalist are going to diagnose Inga and Otto asymmetrically when it comes to whether they count (respectively) as *remembering*—prior to accessing this information from storage—that the Museum of Modern Art is on 53rd street. In Inga's case, we attribute to her a paradigmatic dispositional belief, in virtue of storing this (previously endorsed) information in biomemory. In Otto's case, we—according to the cognitive internalist *and* biointernalist—deny this dispositional belief attribution, simply because the information is not stored in biomemory; it is stored somewhere else.

But how important should *this* be, really? Proponents of extended cognition think that giving this kind of theoretical weight to the material constitution and location of our memory storage is outdated and unprincipled—and as David Chalmers<sup>48</sup> puts it—a form of *bioprejudice*. A more egalitarian approach to the bounds of cognition would have us focus, when assessing whether to include something as part of a cognitive process, less on what it is made of and where it is, and instead on what it does. If something is *doing* the same thing as something that is part of a cognitive process, then why not—in the spirit of parity of treatment—rule it in?

This is the central (then-)radical idea from Clark and Chalmers' (1998) *The Extended Mind*, which is summed up (1998, 8) in their 'parity principle':

Parity Principle: [I]f, as we confront some task, a part of the world functions as a process which, were it to go on in the head, we would have no hesitation in accepting as part of the cognitive process, then that part of the world is part of the cognitive process (1998, 8).

For those willing to reason in accordance with this principle, it will follow that cognition is not *merely* embodied, but also that it can in some circumstances be *extended*—viz. in the sense that not just extracranial but *extraorganismic* things in one's environment (e.g. notebooks, smartphones, tactile vision substitution systems,<sup>49</sup> eyeborgs<sup>50</sup>) can in certain circumstances partially *constitute* that agent's cognitive

<sup>48</sup> See, e.g., Chalmers' forward to Clark's (2008).

<sup>49</sup> (Bach-y-Rita and Kercel 2003). See also Palermos (2014) for discussion.

<sup>50</sup> See Pearlman (2015) for an overview of the case of Neil Harbisson's eyeborg technology, and Carter and Palermos (2016) for discussion of its significance in the wider extended cognition debate.

system. On this way of thinking, we not only can, but should—viz. with reference to the Parity Principle—treat Inga and Otto symmetrically in terms of memory.<sup>51</sup> Not only Inga, but also Otto, remembers—prior to accessing this information from storage—that the Museum of Modern Art is on 53rd street. Additionally, in Inga's case, we attribute a dispositional belief (that the Museum of Modern art is on 53rd street), in virtue of her storing this information in biorecognition. And by parity of reasoning, in Otto's case, we attribute an (extended) dispositional belief with this same content, in virtue of Otto's storing the same information in (extended) memory.

Of course, for the champion of extended cognition, not *everything* that one causally interacts with whilst engaging with a cognitive task is going to get 'ruled in' as part of one's 'extended' cognition. Far from it. One of the key research problems in the contemporary literature on extended cognition is how exactly to distinguish cases like that of Otto, where the parity principle is plausibly satisfied, from cases where we should think it is not—e.g. as when one consults a phone book, or just happens to use a device for a one-off task.<sup>52</sup>

Let us zoom out for a moment. We have been responding in this section to an envisioned critic of the argument from §1. The argument maintained that a cognitive internalist picture of the mind is much more dispensable in epistemology than has been assumed. The anticipated rejoinder was that *even if that is right*, cognitive internalism independently enjoys a kind of sacrosanct status. Granted, the epistemologist (like anyone else) should accept an internalist picture of the mind, and whatever is implied by it, if the cognitive internalism deserves to be treated as a kind of theoretical pillar—one such that we should alter what comes into conflict with it, rather than to alter the pillar itself. *However*, we have now seen in this section that this is hardly the case—and on the contrary—that the tide in recent cognitive science is moving against not only cognitive internalism (e.g. embodied cognition) but also even cognitive biointernalism (e.g. extended cognition).

---

<sup>51</sup> For some recent representative work that discusses different ways to embrace a kind of 'extended cognition' diagnosis of the Otto/Inga case, see, e.g. Clark (2012), Kiverstein (2018), Carter and Kallestrup (2016) and Menary (2010).

<sup>52</sup> As Allen-Hermanson (2013, 793) puts it: 'If a notebook counts as part of one's mind, then why not the yellow pages, the internet, or even parts of the natural world that supply information and support cognition?' For related worries, see Rupert (2004); see also, for discussion, Bjerring and Pedersen (2014) and Pedersen and Bjerring (2021). It is worth registering that the cognitive bloat objection, whilst a common critique of proponents of the extended mind thesis, is best understood as a conditional objection: *if* proponents of the extended mind rule too much in as genuine beliefs (dispositional or otherwise), then this is problematic. Friends of the extended mind thesis accordingly aim to show that their conditions for extending the mind are not too permissive. One strategy, due to Carter and Kallestrup (2020), appeals to 'cluster functionalism', which is meant to be an improvement on Clark's (2010) initial glue and trust conditions. For a more radical proposal, see Palermo (2018), who responds to the problem by denying dispositional beliefs wholesale, via an approach he calls 'epistemic presentism'. Definitely putting the cognitive bloat objection to rest lies beyond the present aim in this section, which has been to call into doubt the sacrosanct status of cognitive internalism, rather than to prove (some version of) its denial. This paper was written as part of the AHRC-funded 'Digital Knowledge: A New Framework for Digital Epistemic Virtues' (AH/W008424/1) project, which is hosted by the University of Glasgow's COGITO Epistemology Research Centre and the University of Aberdeen's CEKAS Centre for Knowledge and Society; I'm grateful to the AHRC for supporting this research.

## 4 Concluding remarks

I have argued here against a certain kind of presumptive picture in epistemology—one that takes cognitive internalism for granted and in doing so de facto forecloses the possibility that epistemic evaluations of good and bad cognition are anything other than intracranial evaluations. This way of thinking artificially—and without satisfactory theoretical motivation (§2)—restricts the way epistemologists approach questions posed by digital information storage and generation, and the truth-directed (epistemic) evaluations we make in response to such questions. When asking whether information stored digitally can be digital *knowledge*, we should feel free to ask this literally, and without concern that doing so requires giving up any underlying commitments we need to make in order to pursue traditional epistemological questions, or (§3) which are otherwise sacrosanct.

## Declarations

**Consent for publication** Confirmed.

**Competing interests** The authors declare no competing interests.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Adams, F. (2010). Embodied cognition. *Phenomenology and the Cognitive Sciences*, 9(4), 619–628.
- Adams, Fred, and Ken Aizawa. (2010). “Defending the bounds of cognition.” In *The extended mind*, edited by Richard Menary, 67–80. MIT Press.
- Adams, Fred, and Kenneth Aizawa. (2009). “Why the mind is still in the head.” In *The Cambridge handbook of situated cognition*, edited by Murat Aydede and P. Robbins, 78–95. Cambridge: Cambridge University Press.
- Adams, Frederick, and Kenneth Aizawa. (2008). *The bounds of cognition*. Wiley-Blackwell.
- Allen-Hermanson, S. (2013). Superdupsizing the mind: Extended cognition and the persistence of cognitive bloat. *Philosophical Studies*, 164(3), 791–806.
- Alston, William P. (1988). “An internalist externalism.” *Synthese*, 265–83.
- Armstrong, D. M. (1969). “Does knowledge entail belief?” *Proceedings of the Aristotelian Society* 70: 21–36. <https://www.jstor.org/stable/4544782>.
- Armstrong, David Malet. (1973). *Belief, truth and knowledge*. CUP Archive.
- Audi, R. (1994). Dispositional beliefs and dispositions to believe. *Noûs*, 28(4), 419–434.
- Bach-y-Rita, P., & Kercel, S. W. (2003). Sensory substitution and the human–machine interface. *Trends in Cognitive Sciences*, 7(12), 541–546. <https://doi.org/10.1016/j.tics.2003.10.013>

- Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72(1–2), 173–215.
- Bernecker, S. (2010). *Memory: A philosophical study*. Oxford University Press.
- Bjerring, J. C., & Pedersen, N. J. L. (2014). All the (many, many) things we know: Extended knowledge. *Philosophical Issues*, 24(1), 24–38.
- Black, C. (1971). Knowledge without belief. *Analysis*, 31(5), 152–158. <https://doi.org/10.2307/3327369>
- Bourget, D., & Chalmers, D. J. (2014). What do philosophers believe? *Philosophical Studies*, 170(3), 465–500. <https://doi.org/10.1007/s11098-013-0259-7>
- Burge, T. (1986). Individualism and psychology. *Philosophical Review*, 95(January), 3–45. <https://doi.org/10.2307/2185131>
- Carter, J. A. (2013). Extended cognition and epistemic luck. *Synthese*, 190(18), 4201–4214.
- Carter, J. A., Gordon, E. C., & OrestisPalermos, S. (2016). Extended emotion. *Philosophical Psychology*, 29(2), 198–217.
- Carter, J. A., & Kallestrup, J. (2016). Extended cognition and propositional memory. *Philosophy and Phenomenological Research*, 92(3), 691–714.
- Carter, J. A., & Kallestrup, J. (2020). Varieties of cognitive integration. *Noûs*, 54(4), 867–890.
- Carter, J. A., JesperKallestrup, S., Palermos, O., & Pritchard, D. (2014). Varieties of externalism. *Philosophical Issues*, 24(1), 63–109.
- Carter, J. A., & OrestisPalermos, S. (2015). Active externalism and epistemic internalism. *Erkenntnis*, 80(4), 753–772.
- Carter, J. A., & OrestisPalermos, S. 2016. “The ethics of extended cognition: Is having your computer compromised a personal assault?”
- Chase, J. (2001). Is externalism about content inconsistent with internalism about justification? *Australasian Journal of Philosophy*, 79(2), 227–246.
- Chemero, A. (2011). *Radical embodied cognitive science*. MIT press.
- Chisholm, R M. (1973). *The problem of the criterion*. Milwaukee.
- Chisholm, R M. (1977). *Theory of knowledge*. Prentice Hall, Englewood Cliffs, NJ,.
- Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. OUP USA.
- Clark, A. 2010. “Memento’s revenge: The extended mind, extended.” *The Extended Mind*, 43–66.
- Clark, A. 2012. “Embodied, embedded, and extended cognition.” *The Cambridge Handbook of Cognitive Science*, 275–91.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19. <https://www.jstor.org/stable/3328150>
- Conee, E. (2004). *and Richard Feldman*. Essays in epistemology. Clarendon Press.
- Cottingham, J. (2002). Descartes and the voluntariness of belief. *The Monist*, 85(3), 343–360.
- Craig, E. (1991). *Knowledge and the state of nature: An essay in conceptual synthesis*. Clarendon Press.
- Cummins, R. (1996). *Representations, targets and attitudes*. MIT Press.
- Brigard, De., & Felipe. (2014). The nature of memory traces. *Philosophy Compass*, 9(6), 402–414.
- Dutant, J. (2015). The legend of the justified true belief analysis. *Philosophical Perspectives*, 29(1), 95–145.
- Farkas, K. (2015). Belief may not be a necessary condition for knowledge. *Erkenntnis*, 80(1), 185–200.
- Fodor, J A., and Fodor, J A. (1987). *Psychosemantics: The problem of meaning in the philosophy of mind*. Vol. 2. MIT press.
- Gallagher, S. (2014). “Phenomenology and embodied cognition.” *The Routledge Handbook of Embodied Cognition*, 9–18.
- Gerken, M. (2015). “How to do things with knowledge ascriptions.” Edited by Krista Lawlor. *Philosophy and Phenomenological Research* 90 (1): 223–34. <https://www.jstor.org/stable/24672859>.
- Gettier, E L. 1963. “Is justified true belief knowledge?” *Analysis* 23 (6). <https://doi.org/10.1093/analys/23.6.121>.
- Gibbs, R. W., Jr. (2005). *Embodiment and cognitive science*. Cambridge University Press.
- Gibbs, R. W., Jr. (2005). *Embodiment and cognitive science*. Cambridge University Press.
- Goldman, A. I. (1979). “What is justified belief?” In *Justification and knowledge* (pp. 1–23). Springer.
- Haidt, J., Koller, S. H., & Dias, M. G. (1993). Affect, culture, and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology*, 65(4), 613.
- Hannon, M. (2018). *What’s the point of knowledge?: A function-first epistemology*. Oxford University Press.
- Huemer, M. (2006). Phenomenal conservatism and the internalist intuition. *American Philosophical Quarterly*, 43(2), 147–158.

- Hurlley, S. L. (1998). *Consciousness in action*. Harvard University Press.
- Ichikawa, J. J., & Matthias S. (2018). "The analysis of knowledge." In *The Stanford encyclopedia of philosophy*, edited by Edward N. Zalta, Summer 2018. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/sum2018/entries/knowledge-analysis/>.
- Kiverstein, J. (2018). "Extended cognition." *The Oxford Handbook of 4e Cognition*, Oxford University Press, New York, 19–40.
- Kriegel, U. (2013). *Phenomenal intentionality*. Oxford University Press.
- Krueger, J., & Szanto, T. (2016). Extended emotions. *Philosophy Compass*, 11(12), 863–878. <https://doi.org/10.1111/phc3.12390>
- Lakoff, G. (2012). Explaining embodied cognition results. *Topics in Cognitive Science*, 4(4), 773–85. <https://doi.org/10.1111/j.1756-8765.2012.01222.x>
- Lehrer, K. (1968). Belief and knowledge. *The Philosophical Review*, 77(4), 491–499.
- Lehrer, K. (1990). 2018. Routledge.
- Loar, B.R. Grimm, D. Merrill. (1988). "Social content and psychological content." *Contents of Thought*.
- Lycan, William. (2000). "Representational theories of consciousness." May. [https://plato.stanford.edu/entries/consciousness-representational/?utm\\_content=buffer1d085&utm\\_medium=social&utm\\_source=pinterest.com&utm\\_campaign=buffer](https://plato.stanford.edu/entries/consciousness-representational/?utm_content=buffer1d085&utm_medium=social&utm_source=pinterest.com&utm_campaign=buffer).
- Lycan, W G., (1986). "Tacit belief." In *Belief: Form, content, and function*, edited by R. Bogdan. Oxford University Press.
- Madison, Brent J.C.(2017). "Internalism and externalism." *The Routledge Handbook of Philosophy of Memory*, 283.
- Markie, P., (1992). "The cogito and its importance." In *Descartes*, edited by John Cottingham. Oxford University Press.
- McBeath, M. K., Shaffer, D. M., & Kaiser, M. K. (1995). How baseball outfielders determine where to run to catch fly balls. *Science*, 268(5210), 569–573.
- McKenna, Robin. (2013). "'Knowledge' ascriptions, social roles and semantics." 2013. <https://philarchive.org>.
- McNeill, D. (1992). *Hand and mind: What gestures reveal about thought*. Univ. Chicago Press.
- Menary, R. (2010). *The extended mind*. Mit Press.
- Miracchi, L., (2021). "Embodied cognition and the causal roles of the mental." In *Mental action*, edited by Michael Brent. London: Routledge.
- Moore, G. E. (1962). *Commonplace book: 1919–1953*. Routledge.
- Murray, D., Sytsma, J., & Livengood, J. (2013). God knows (but does God believe?). *Philosophical Studies*, 166(1), 83–107. <https://doi.org/10.1007/s11098-012-0022-5>
- Myers-Schulz, B., & Schwitzgebel, E. (2013). Knowing that P without believing that P. *Noûs*, 47(2), 371–384.
- Noë, A. (2005). *Action in perception*. MIT Press.
- Palermos, S. O. (2014). Loops, constitution, and cognitive extension. *Cognitive Systems Research*, 27, 25–41.
- Palermos, S. O. (2016). The dynamics of group cognition. *Minds and Machines*, 26(4), 409–440.
- Palermos, S. O. (2018). Epistemic presentism. *Philosophical Psychology*, 31(3), 458–478. <https://doi.org/10.1080/09515089.2018.1432035>
- Palermos, S. O., & Pritchard, D. (2013). Extended knowledge and social epistemology. *Social Epistemology Review and Reply Collective*, 2(8), 105–120.
- Pearlman, E., (2015). "I, cyborg." *PAJ. A Journal of Performance and Art* 37. (2): 84–90
- Pedersen & Bjerring (2021). 'Extended knowledge overextended?' in K. Lai (ed.): *Extending knowledge: Reflections on epistemic agency and epistemic environment in East-West philosophy* (Palgrave Macmillan).
- Poston, T. (2014). "Internalism and externalism." *Internet Encyclopedia of Philosophy*.
- Pritchard, D. (2010). Cognitive ability and the extended cognition thesis. *Synthese*, 175(1), 133–151. <https://doi.org/10.1007/s11229-010-9738-y>
- Pritchard, D., & Kallestrup, J. (2004). An argument for the inconsistency of content externalism and epistemic internalism. *Philosophia*, 31(3), 345–354.
- Putnam, H. (1975). The meaning of 'meaning.' *Minnesota Studies in the Philosophy of Science*, 7, 131–193.
- Radford, C. (1966). Knowledge: By examples. *Analysis*, 27(1), 1–11. <https://doi.org/10.2307/3326979>
- Ramsey, W. (2017). Must cognition be representational? *Synthese*, 194(11), 4197–4214.

- Rose, D., & Schaffer, J. (2013). Knowledge entails dispositional belief. *Philosophical Studies*, 166(1), 19–50.
- Rupert, R. D. (2004). Challenges to the hypothesis of extended cognition. *The Journal of Philosophy*, 101(8), 389–428.
- Schellenberg, S. (2017a). “Perceptual capacities, knowledge, and Gettier cases.” In *Explaining knowledge: New essays on the Gettier problem*, 74–95. Oxford: Oxford University Press.
- Schellenberg, S. (2017b). “The origins of perceptual knowledge.” *Episteme* 14 (3): 311–28. <https://doi.org/10.1017/epi.2017b.22>.
- Schwitzgebel, E. (2019). “Belief.” In *The Stanford encyclopedia of philosophy*, edited by Edward N. Zalta, Fall 2019. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2019/entries/belief/>.
- Setiya, K. (2008). Practical knowledge. *Ethics*, 118(3), 388–409.
- Shapiro, L. (2014). *The Routledge handbook of embodied cognition*. Routledge.
- Shea, N. (2013). Naturalising representational content. *Philosophy Compass*, 8(5), 496–509.
- Shope, Robert K. (2017). “The analysis of knowing: A decade of research.”
- Simon, D. T., Larsson, K. C., Nilsson, D., Burström, G., Galter, D., Berggren, M., & Richter-Dahlfors, A. (2015). An organic electronic biomimetic neuron enables auto-regulated neuromodulation. *Biosensors and Bioelectronics*, 71, 359–364.
- Sosa, E. (1980). The raft and the pyramid: Coherence versus foundations in the theory of knowledge. *Midwest Studies in Philosophy*, 5, 3–25.
- Sosa, E. (1997). “Reflective knowledge in the best circles.” *Journal of Philosophy* 94 (8): 410–30. [jphil199794827](http://phil199794827)
- Sosa, E. (2015). *Judgment & agency*. Oxford University Press UK.
- Stalnaker, R. (1973). Presuppositions. *Journal of Philosophical Logic*, 2(4), 447–457.
- Sutton, J. (2006). Introduction: Memory, embodied cognition, and the extended mind. *Philosophical Psychology*, 19(3), 281–289.
- Vetere, G., Tran, L. M., Moberg, S., Steadman, P. E., Restivo, L., Morrison, F. G., Ressler, K. J., Josselyn, S. A., & Frankland, P. W. (2019). Memory formation in the absence of experience. *Nature Neuroscience*, 22(6), 933–940.
- Wheeler, M. (2018). *Knowledge, credit, and the extended mind, or what Calvisius Sabinus got right*. Oxford University Press. <http://dspace.stir.ac.uk/handle/1893/26500>.
- Williamson, T. (2002). *Knowledge and its limits*. Oxford University Press on Demand.
- Wilson, R A., and Lucia F. (2017). “Embodied cognition.” In *The Stanford encyclopedia of philosophy*, edited by Edward N. Zalta, Spring 2017. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2017/entries/embodied-cognition/>.
- Yli-Vakkuri, J., & Hawthorne, J. (2018). *Narrow content*. Oxford University Press.

The content of the manuscript has not been published, or submitted for publication elsewhere (see our Duplicate publication policy).