

# Extended cognition and epistemic luck

J. Adam Carter

Received: 26 December 2012 / Accepted: 27 February 2013 / Published online: 30 March 2013  
© Springer Science+Business Media Dordrecht 2013

**Abstract** When extended cognition is extended into mainstream epistemology, an awkward tension arises when considering cases of environmental epistemic luck. Surprisingly, it is not at all clear how the mainstream verdict that agents lack knowledge in cases of environmental luck can be reconciled with principles central to extended cognition.

**Keywords** Epistemic luck · Extended cognition · Knowledge

## 1 Introduction

Proponents of extended cognition lament what [Clark \(2011\)](#) calls *bioprejudice*. If, for some process  $C$ , one would count  $C$  as a cognitive process were it to have occurred within the skull and skin, *but exclude it as such because it does not*, the kind of alleged bioprejudice here would be metaphysical—a prejudice that affects one’s position on the metaphysical nature of a cognitive process. [Clark and Chalmers \(1998\)](#), in their seminal paper, think we can avoid the mistake of metaphysical bioprejudice by adhering to what they call the *parity principle*:

PARITY PRINCIPLE: If, 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. ([Clark and Chalmers 1998](#), p. 8)

---

J. A. Carter (✉)  
Edinburgh University, Edinburgh, UK  
e-mail: jadamcarter@gmail.com

J. A. Carter  
Eindhoven University of Technology, Eindhoven, The Netherlands

Those who accept the parity principle will count as cognitive processes the *extracranial analogues* of whatever intracranial cognitive processes they already recognize. For some intracranial cognitive process  $C$ ,  $C^*$  is  $C$ 's extracranial<sup>1</sup> analogue just in case, by reference to the parity principle,  $C^*$  is a cognitive process if  $C$  is a cognitive process. In Clark and Chalmers's classic case of Otto, the process of consulting a notebook is presented as the extracranial analogue of consulting one's biological memory.

*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.

The parity principle forbids metaphysical bioprejudice while permitting what I'll call *epistemic bioprejudice*. Consider that I could, by reference to the parity principle, recognize Otto's process as a *bona fide* cognitive process, whilst *nonetheless* maintaining the following: that Otto fails to count as knowing<sup>2</sup> some proposition  $p$ , even though Otto's intracranial counterpart, who differs from Otto only in that he relies on a biological memory, counts as knowing  $p$ <sup>3</sup>.

To the extent that proponents of extended cognition appeal to egalitarian considerations, they should be dissatisfied by the prospect that the extracranial nature of Otto's cognitive process could ever exclude him from knowing what his counterpart, different only in relying on the intracranial analogue of the process Otto uses, counts as knowing. More than the parity principle is needed, then, to combat bioprejudice in epistemology.

## 2 Epistemic parity

The parity principle then guards against only *metaphysical* bioprejudice. That's not good enough. Proponents of cognitive extension require an additional principle, an *epistemic* parity principle, to guard against epistemic bioprejudice. Call this the *E-parity principle*.

<sup>1</sup> I am using this term to refer to processes claimed to extend, at least in part, outside the skin and skull. I am not making a distinction here between extracranial and what some writers have called *transcranial* processes. See here, for example, Adams and Aizawa (2008).

<sup>2</sup> For simplicity's sake, I am framing epistemic bioprejudice in terms of knowledge: one betrays her epistemic bioprejudice when withholding knowledge in a case where the process employed is in part extracranial and, were the process employed the *intracranial analogue* of the process actually employed, knowledge would have been ascribed. There are of course other ways to betray epistemic bioprejudice, insofar as cases of extended cognition stand subject to our ascriptions of (for example) justification, understanding, rationality or intellectual virtue.

<sup>3</sup> Such a verdict would betray epistemic, rather than metaphysical, bioprejudice because the stance welcomes extracranial processes as cognitive processes; the bioprejudice is epistemic because it takes the distinction to be an epistemic difference maker, even if not a difference maker from the perspective of assessing what qualifies as a cognitive process.

**E-PARITY PRINCIPLE:** For agent  $S$  and belief  $p$ , if  $S$  comes to believe  $p$  by a process which, were it to go on in the head, we would have no hesitation in ascribing knowledge of  $p$  to  $S$ , then  $S$  knows  $p$ .

Just as the original parity principle implies that we must accept as cognitive processes the *extracranial analogues* of whatever intracranial cognitive processes we already recognize, the E-parity principle implies something to the effect that: we must accept as cases of knowledge the *extracranial epistemic analogue cases* of whatever intracranial cases of cognition we count as knowledge. For two otherwise epistemically symmetrical<sup>4</sup> cases of cognition,  $A$  and  $A^*$ ,  $A^*$  is an extracranial epistemic analogue (hereafter, extracranial analogue) of  $A$  just in case,  $A$  and  $A^*$  differ in that the process employed in  $A^*$  is, by reference to the (original) parity principle, the extracranial analogue of the cognitive process employed in  $A$ . More generally, the E-parity principle implies that substituting the extracranial analogue of an intracranial cognitive process, while keeping everything else epistemically symmetrical across cases, should never be difference-making from the perspective of epistemic evaluation.<sup>5</sup> Even shorter: the difference between intracranial and extracranial analogues is not an epistemic difference.

The E-parity principle captures I think a key implication of extended cognition for epistemological theorising<sup>6</sup>: it constitutes, no less than the parity principle, a kind of *ex ante* constraint on any epistemological thesis amenable to cognitive extension. Any epistemological thesis violates the E-parity principle if it has as a consequence that (for instance) Otto's intracranial analogue, but not Otto, counts as a knower, just as any epistemological thesis violates the parity principle if it counts Otto's intracranial analogue but not Otto as employing a cognitive process. Commitments in epistemology can stand in tension with *either* the parity principle or the E-parity principle, and the egalitarian aims of cognitive extension should permit the violation of neither.

### 3 Parity and epistemic luck

Epistemologists overwhelmingly accept that the following is *not* a case of knowledge<sup>7</sup>:

<sup>4</sup> If two cases are epistemically symmetrical, then (for instance) a theory of epistemic justification or knowledge must treat them the same.

<sup>5</sup> For two cases,  $A$  and  $A^*$ , where  $A^*$  is the extracranial analogue case of  $A$ , non-epistemically relevant features of  $A^*$  can differ from  $A$  so long as the *epistemically relevant* features of the case—*viz.*, the features that matter for assessing whether the case is a case of knowledge—remain symmetrical. Because E-parity is a principle aimed at preventing prejudice about what counts as knowledge, it is appropriate to frame E-parity in terms of analogue cases that require just that the relevant epistemic features are held fixed (rather than in terms of the more narrow category of intercranial–extracranial *counterpart* cases, according to which even the non-epistemically relevant features across cases are held fixed.)

<sup>6</sup> See here especially Hetherington (2012) *gradualist* proposal, according to which, with reference to extended cognition, knowledge might be attributed to different extents.

<sup>7</sup> For a sample of some explicit statements on this point, see for example Swain (1978), Lewis (1996), Kvanvig (2003, 2004), Pritchard (2003), (2004), (2005), (2007), (2012), Steup (2008), Luper (2010), Madison (2011), Kelp (2012a), Jarvis (2012), Carter (2010, 2011) and Carter et al. (2012a,b). This list is far from exhaustive. There is some dissent, however—especially from experimental philosophy—which I consider in §4.

**Barn facade:** Driving through the countryside, Jimmy wants to teach his son the word *barn*. He points to what (unbeknownst to him) is the lone real barn in the area. He could have just as easily pointed to one of the facades (indistinguishable from the real barn) and formed a false belief.

The (very) short explanation for why Jimmy lacks knowledge is that his getting it right is down to luck. Holding fixed the way Jimmy forms the belief in the actual world—by pointing to one of the barn-looking objects in the facade-littered countryside—he believes falsely in nearby possible worlds. Compare the barn facade case now with the following case in which we reprise Otto from the original [Clark and Chalmers \(1998\)](#) example:

**Case 1:** Otto consults his notebook to determine when his doctor’s appointment was today, and finds the correct time, noon, written in the book. Unbeknownst to Otto, his notebook had been stolen by a jokester, who fudged with the times of Otto’s other appointments that day, changing them all back an hour. The jokester, however, overlooked the doctor’s appointment, leaving the original and correct time intact.

A natural response to Case 1 is to view it as structurally akin to the barn facade case. We should deny knowledge to Otto for the very same reasons we deny knowledge to Jimmy<sup>8</sup>. As epistemologists typically put it, both cases are ones where knowledge is undermined by *environmental* epistemic luck. On this more will be said. But first, some more general observations: that Otto fails to count as knowing in Case 1 is not itself problematic. But this verdict—perfectly intuitive in mainstream epistemology—has two very awkward implications for proponents of extended cognition. The first awkward implication is that it stands in tension with the parity principle. The second awkward implication is that it stands in tension with the E-parity principle. Let’s look at these in turn.

### 3.1 Tension with (metaphysical) parity

That the no-knowledge verdict given to Case 1 stands at tension with metaphysical parity (hereafter, *M-parity*) is not obvious. But after a few background remarks about epistemic luck, it will be.

Consider first that a task of *anti-luck epistemologists* is to specify the precise sense in which knowledge cannot depend on luck<sup>9</sup>, and to generate (with reference to this specification), a plausible *anti-luck condition* on knowledge. The anti-luck condition

<sup>8</sup> Consider that, while one who endorses extended cognition will, by reference to the parity principle, count the notebook employed in the actual world as part of the cognitive process, the same is not true for the barn facades. Note that a condition for satisfying the parity principle (which circumscribes which processes are counted as cognitive processes, by proponents of extended cognition), involves Clark and Chalmers’s (1998) notion of *coupling*. While Otto counts as coupled to his notebook, Jimmy does not qualify as coupled to the genuine barn. See also §5 for additional discussion here. Thanks to an anonymous reviewer for raising this issue.

<sup>9</sup> As Pritchard (2007) puts it: “...we evaluate theories of knowledge in terms of whether they are able to accommodate this claim. If they can’t—that is, if they allow lucky knowledge—then this is taken to be a decisive ground for rejecting the view.”

needs to not be so strong that it gives false positives. After all, sometimes, knowledge can depend on luck in a way that is epistemically *benign*. Consider the following.

1. My knowledge that the Swedes, despite being outnumbered 10-to-1, defeated the Russians in the battle of Narva in 1700.
2. Detective Fusco’s knowledge that Johnny was planning a murder—knowledge Fusco acquires by luckily finding Johnny’s confession note.

The first case is a case of *content epistemic luck*. What is lucky here is just *that* the proposition is true. The Swedes could have easily lost that battle (and almost surely would have, had there not been a blinding snowstorm).<sup>10</sup> Knowledge is perfectly compatible with content epistemic luck; we know of many improbable events, that they occurred. The luck present in (2) is different, though equally unproblematic from the perspective of ascribing knowledge. In (2), what is lucky is that Fusco ever acquired the evidence that he does in favour of his belief. Call this kind of luck *evidential epistemic luck*. The fact that one’s being correct depends on evidential epistemic luck is not something our ascriptions of knowledge are, or should be, sensitive to. So what kind of luck can knowledge *not* depend on? The dominant answer in the literature is Pritchard’s (2005): knowledge is incompatible with what is called *veritic epistemic luck*<sup>11</sup>:

VERITIC EPISTEMIC LUCK: For all agents *S* and propositions *p*, the truth of *S*’s belief that *p* is veritically lucky if and only if *S*’s belief that *p* is true in the actual world *a* but false in nearly all nearby possible worlds in which *S* forms the belief in the same manner as in *a*.

Veritic luck comes in two notable varieties: *intervening veritic luck* and *environmental veritic luck*. The former variety is characteristic of standard Gettier-style cases, where luck *intervenes* between the agent the belief she forms<sup>12</sup>. In cases of environmental veritic luck, that you could easily have been wrong is down to a bad epistemic environment<sup>13</sup> rather than to any luck that interferes with the agent’s formation of a true belief<sup>14</sup>.

The claim that knowledge can’t depend on veritic luck (of either variety) is tantamount to saying that: if you know, then you *could not easily have been wrong* (given the way you came to form the target belief). For the present purposes, what’s important to note is that an assessment of whether the correctness of a given belief depends on

<sup>10</sup> I borrow this example from Peterson (2009, p. 4).

<sup>11</sup> For a recent challenge to this position, see Hetherington (2011), Ch. 3; here Hetherington defends what he calls combinatorial luck as better-suited than veritic luck at capturing the luck that is incompatible with knowing. Hetherington (forthcoming) also challenges Pritchard on this point; compare here with Pritchard’s (forthcoming) reply, in the same volume.

<sup>12</sup> Consider, for example, Gettier’s (1963) original case involving Smith, who believes, (with justification), that the man who gets the job has 10 coins in his pocket. Smith’s belief is true, but true because a different man (himself) happened to get the job, and Smith also happened to have 10 coins in his pocket. Cf. §4.2, Case 5.

<sup>13</sup> Thanks to an anonymous reviewer for pointing out that, often, bad epistemic environments are presented as ones that are also unfamiliar to the agent. It should be stressed that unfamiliarity is (though common) not an *essential* feature of a bad epistemic environment. See here Zagzebski’s (1996), pp. 285–286) environmental-luck style challenge to Plantinga’s (1993).

<sup>14</sup> See Pritchard (2007) for further discussion on this point.

veritic epistemic luck (hereafter, epistemic luck) requires a kind of *framing problem*: we have to decide what to hold fixed under the description of the relevant manner that led to the formation of the target belief in the actual world, when we move out to nearby worlds. The framing problem is a species of the generality problem<sup>15</sup>: there are costs associated with construing what to hold fixed too narrowly, as well as too widely. Accordingly, it's hard to specify precisely *all* that needs to be held fixed<sup>16</sup>. But we will always *at least* hold fixed everything under the description of the *cognitive process* the agent employs in the actual world, as well as such things as the agent's physical and mental condition (Pritchard 2007). And what to hold fixed under the description of the cognitive process deployed will be partially, though not wholly, determined by one's metaphysical commitments as to the nature of cognitive processes<sup>17</sup>.

We are now in a position to see how the no-knowledge verdict in Case 1 threatens M-parity. Otto's belief is epistemically lucky because, given how he forms the belief—by consulting a notebook tampered with by a jokester—he could have easily been incorrect—e.g. in many close possible worlds, the jokester *does not* overlook the doctor's appointment entry, and Otto is an hour early to the appointment. The key point here is that, by reference to M-parity, *the notebook itself*—a notebook with a *correct* time written in it—is held fixed under the description of the relevant cognitive process. But if it is, then the safety of Otto's belief is no longer undermined by environmental epistemic luck in Case 1<sup>18</sup>. After all, in most nearby worlds where Otto consults a notebook with a correct time for the doctor appointment, he forms a true belief—and so his correctness isn't lucky<sup>19</sup>. (Compare: if we held fixed the barn Jimmy actually points to in the barn facade case, we'd get the result that he knows—the wrong result). The upshot here is that the verdict that Otto lacks knowledge in Case 1 seems to fly squarely in the face of M-parity. M-parity implies that Otto's process—the extracranial analogue of a cognitive process (memory) we recognize—is a cognitive process. But we most saliently explain why Otto's belief is unsafe (and therefore, not knowledge<sup>20</sup>)

<sup>15</sup> Thanks to an anonymous reviewer for raising this point.

<sup>16</sup> Cf. Kvanvig (2008), in his critique of Pritchard (2005), for some observations on the difficulty of getting the framing problem right. See also Pritchard (2010).

<sup>17</sup> Accordingly, proponents of extended cognition who are metaphysically committed to allowing more of the world to feature in a description of a cognitive process will hold more fixed under the description of the cognitive process employed in the actual world, than those without such commitments.

<sup>18</sup> The same results follows if we think of safety as a property of methods fundamentally, and beliefs derivatively. The extended memorial method employed by Otto in the actual world is a method of consulting his notebook, as it is in the actual world. This method is safe in nearby worlds, as the notebook contains the correct entry. Thanks to Chris Kelp for discussion on this point.

<sup>19</sup> Indeed, if we hold fixed that Otto consults a notebook with the *correct* time of his doctor's appointment, the nearest worlds where he gets the time wrong are far-off worlds—worlds where (for instance) something unusual (say, bizarre ambient light) prevents him from apprehending the time of his appointment when he consults a notebook in which it is clearly written.

<sup>20</sup> Although I think that safe, true beliefs also happen to often satisfy the conditions for knowledge, I am not committing to safety as a sufficient condition for knowledge. Thanks to a reviewer for pointing out that, in any of the cases discussed here, there are features other than safety which could potentially be referenced in accounting for why an agent does not know. What's relevant to my discussion on this point is just that, given M-Parity, the salient explanation for why Otto fails to know—namely, that is belief is unsafe—is off the table.

by *excluding the notebook* (which has the correct time) from what we hold fixed under the description of the cognitive process Otto employs in forming a belief about the time of his doctor's appointment. And such an exclusion is in violation of M-Parity.

### 3.2 Tension with E-parity

Let's revisit Case 1, which by the lights of mainstream epistemology is not a case of knowledge, but rather, a case in which knowledge is undermined by environmental epistemic luck. By reference to the E-parity principle, knowledge should also fail to be present in any case that is an *intracranial analogue case* of Case 1. To get an intracranial analogue case for Case 1, we first substitute biological memory for the notebook. Next, we keep the epistemically relevant features fixed across cases; this involves—in the case of a biological memory—obscuring the other appointment times recorded in Otto's memory, though permitting him to remember his doctor's appointment time<sup>21</sup>. Consider now, Case 2, which meets these conditions for an intracranial analogue case to Case 1:

**Case 2:** Otto\* (without Alzheimers) has a normally functioning biological memory, which he relies on to organise his world. Atypically for Otto\*, he forgets the time of his other appointments today—believing they were earlier than they actually are—though he does remember that his doctor's appointment is at noon.

Since Case 1 is an extracranial analogue of Case 2, E-parity tells us that the Case 1 is a case of knowledge if Case 2 is<sup>22</sup>. Now here's the twist: Mainstream thinking about epistemic luck tells us Case 2 *is* a case of knowledge<sup>23</sup>. Given the way Otto\* forms his belief—*viz.*, *by consulting a clear memory of the time of the appointment*—he could not easily have been wrong. Does it matter that, even though he didn't, he could

<sup>21</sup> Thanks to Benjamin Jarvis, Chris Kelp and Wybo Houkes for helpful correspondence on this point. We can think of the role of the jokester in this analogue as one's neurological makeup, the latter of which determines that some, but not other, memories are obscured.

<sup>22</sup> Thanks to an anonymous reviewer for pointing out that one might object to the analogy by arguing as follows: that there is a factual defeater present in Case 1 (namely, that the jokester tampered with the other dates in the notebook) but not in Case 2. However, on closer inspection, the same kind of factual defeater either is or is not present in both cases. Consider that in Case 2, one could have forgotten the target proposition (by the same mechanisms that led to Otto\* forgetting the other appointments). Deny that this fact is a factual defeater, and it's hard to see how to motivate the line that there is a factual defeater in Case 1 consisting in the fact that the date *could* have easily been changed by the jokester, even though it wasn't. I should add that, as an aside, I don't want to take a stance on the epistemic status of factual defeaters, or the conditions for factual defeat. The point is only that the cases are more similar with respect to apparent factual defeat than it might seem on first blush.

<sup>23</sup> Hetherington (1998, forthcoming) has suggested that cases like Case 2 count against the anti-luck platitude, as they constitute a kind of lucky knowledge. He goes even further to conclude that knowledge is present in Gettier cases. For a recent criticism of Hetherington on these points, see Madison (2011). It is worth pointing out that anti-luck epistemologists can agree with Hetherington that Case 2 is a case of lucky knowledge while denying that the case is lucky in a way that is, by reference to veritic luck, incompatible with knowledge. Put differently, the sense in which we recognise Case 2 as "lucky knowledge" does not show knowledge to be compatible with veritic luck anymore than recognising "fool's gold" commits one to the view that there is a kind of gold coextensive with iron pyrite.



have forgotten his doctor's appointment time, as he'd forgotten the times of his other appointments? Consider Case 3:

**Case 3:** Phineas is working on the railroad and slips, coming within an inch of impaling his frontal lobe with an iron rod. Had the iron rod impaled his frontal lobe, he would have confused the time of his doctor's appointment (by thinking it was an hour earlier). Fortunately for Phineas, the iron rod merely affected the part of his brain responsible for mood (making him a bit angrier), and after he stands up, he remembers correctly the time of his doctor's appointment.

In Case 3, the natural verdict is that Phineas knows the time of his doctor's appointment. And that's indeed the right result. Although Phineas surely could have damaged his memory, he didn't, and as he didn't it's not lucky that he's correct. Likewise, in Case 2, Otto\* (like anyone else) could have forgotten the time of his appointment, though he didn't, and because he clearly remembers it, it's not lucky that (in retrieving what he clearly remembers) he's correct<sup>24</sup>.

But if Case 1 is not a case of knowledge and Case 2 is, then we have a conflict with E-parity. E-parity tells us that Case 1 is a case of knowledge if Case 2 is. So it looks like we have to either count Case 1 as a case of knowledge (and fly in the face of anti-luck epistemology) or give up on E-parity (and fly in the face of extended cognition). Framed this way, our ordinary thinking about epistemic luck commits bioprejudice, or, proponents of extended cognition should restrict their egalitarian ambitions to metaphysical parity, giving up on epistemic parity.

#### 4 Responses

As things stand, it looks like mainstream thinking about epistemic luck cannot easily make room for extended cognition, any more than extended cognition can make sense of mainstream thinking about epistemic luck<sup>25</sup>. So far, I've argued that the mainstream verdict in Case 1 flies in the face of both the original parity principle as well as the E-parity principle (a principle I argued in §2 is needed to guard against bioprejudice in epistemology). But perhaps I've been too quick. Let's investigate now how we might escape this unhappy conclusion.

All the tension evaporates, of course, if Otto simply counts as knowing in Case 1<sup>26</sup>. The similarity between Case 1 and a barn facade case is difficult to deny. So trying to

<sup>24</sup> After all, the nearest worlds in which he gets it wrong despite retrieving a correct memory on the point are worlds where there is some anomalous, neural breakdown that occurs *after* the retrieval of the correct memory, blocking the formation of a correct belief.

<sup>25</sup> Compare this assessment with Plantinga's (1993) point that Gettier cases depend on JTB's allowing its justification condition to be satisfiable entirely by an internalistic justification condition. Thanks to an anonymous referee for drawing attention to this comparison.

<sup>26</sup> Why not simply say that Jimmy has a veridical perceptual experience in the barn facade case, and leave it at that? This is the broadly McDowell (1994) line. And since Jimmy satisfies McDowell's further condition of not believing irresponsibly, Jimmy knows, on McDowell's view. I concede this point, though I note that the cost of the McDowellian line is an (apparent) inability to handle cases where environmental epistemic luck is widely thought to undermine knowledge. Thanks to an anonymous reviewer for noting the availability of this line of reply.



save Case 1 as a case of knowledge by granting that barn facade cases are not cases of knowledge, and then arguing Case 1 to be structurally dissimilar to a barn facade case, is unpromising<sup>27</sup>. Another strategy would be to grant the structural similarity between Case 1 and barn facade cases while parting with orthodoxy and claiming these cases actually *are* cases of knowledge. Such a strategy has some precedent, though not much.

This line of thought is developed, in a qualified way, by Sosa (2007), who distinguishes between *animal knowledge* and *reflective knowledge*<sup>28</sup>. Sosa thinks reflective knowledge is not compatible with the kind of environmental luck at play in barn facade cases, though he is willing to grant (lower-grade) animal knowledge in such cases. One could motivate an ascription of knowledge in Case 1, then, by arguing that animal knowledge is present, even if reflective knowledge is not. One problem with this route, though, is that extended cognition would be (in a qualified way) reconcilable with mainstream thinking about epistemic luck only by one who already grants Sosa's unorthodox distinction between animal and reflective knowledge. While this allows for a reconciliation, it does so in a way that is problematically restrictive, pushing the problem down the line to the majority of epistemologists who do not subscribe to a dual account of animal and reflective knowledge.

Another kind of strategy would be to appeal to experimental philosophy, or X-phi, in an attempt to challenge the broad consensus in epistemology regarding cases of environmental luck. In a recent paper on epistemic intuitions and barn facade thought experiments, Colao et al. (2012) found that, while individuals surveyed did distinguish between barn facade and control cases (where the agent is in a good environment), the difference is not as significant as would be suggested by the overwhelming consensus in epistemology. They also found that *age* matters; older individuals surveyed were less likely to attribute knowledge than were younger people surveyed. Unfortunately, it's not clear what surveys like this are supposed to mean for epistemologists. This is a large issue, and not one I will attempt to address here. It should suffice to point out that the kind of resolution we should want between mainstream anti-luck epistemology and extended cognition should not be a resolution that can be purchased only by those willing to explain how X-phi data is supposed to overturn mainstream philosophical consensus<sup>29</sup>.

If we are going to resolve the tension between the mainstream verdict on Case 1 and the M-Parity and E-Parity principles, it won't be as easy as rejecting the mainstream verdict on Case 1. Something else has to give. What has to give does not necessarily have to be E-Parity or M-Parity. Perhaps the problem lies with some of the other claims we relied on to generate the tension between the Case 1 and M-Parity and E-Parity, respectively.

<sup>27</sup> Cf. Kelp (2012b) for a related discussion on this point.

<sup>28</sup> On Sosa's view, animal knowledge is apt belief—belief that is accurate *because* adroit. Reflective knowledge is meta-apt belief, which is apt belief, aptly noted. Unlike animal knowledge, reflective knowledge requires that the agent exhibit a kind of meta-competence, whereby she recognises (or takes for granted) that her first-order belief as aptly formed.

<sup>29</sup> See Knobe and Nichols (2008) for a sample of papers supporting and criticising the experimental approach.

#### 4.1 The tension with M-Parity (revisited)

The tension between M-Parity and a recognition of Case 1 as a case of knowledge was, in short, that M-Parity demands we include Otto's notebook as part of his cognitive process. But since the notebook had the *correct* time for his doctor's appointment in it, most nearby worlds in which he consults *that* notebook are ones where he gets it right. Since that notebook is, by reference to M-Parity, part of his cognitive process, it comes with him when we move out from the actual world. But because it does, it seems that an implication of M-Parity is that Otto's belief is safe, as he believes truly in most nearby worlds. But of course, this result conflicts with the verdict that his belief is undermined by environmental epistemic luck.

We are not jettisoning the verdict that Otto lacks knowledge, and neither should we be comfortable with so quickly dismissing the independently plausible M-Parity. So let's look at the details more closely to determine what else is on the table. As we considered in §3.1, we can make sense of Otto's belief (in Case 1) as undermined by environmental luck if we narrow the cognitive process to exclude the notebook. But this simply conflicts with M-Parity. What if, instead, we beefed up the cognitive process, to include not only the notebook but also the *jokester*? Such a proposal has an initial promise for resolving the conflict: we might think that, whilst holding fixed the notebook (with the correct appointment time) was responsible for the result that Otto's belief would be safe, we could counteract this result by including the tampering jokester under the description of the cognitive process. It might be argued, then, that since the jokester is tampering with the notebook in nearby worlds, Otto's belief is not knowledge (consistent with the mainstream verdict in epistemology), *and* since we are not excluding the notebook from the cognitive process, we are not at odds with M-Parity.

Unfortunately, this line runs into some disastrous problems. Firstly, if we take seriously the idea of holding fixed the jokester (as well as the notebook) as part of the cognitive process, then we are holding fixed both the jokester and notebook from the actual world. This means holding fixed a notebook with the correct time for the doctor's appointment *and* a *jokester who overlooked that*. But in most nearby worlds in which Otto forms his belief by consulting a notebook with a correct time (which a tag-a-long jokester overlooks when changing the others), Otto believes truly, and thus, the belief is not undermined by environmental epistemic luck. This problem is devastating. But even if it could be avoided, there is the separate problem, which is that it's not clear that, even on the most liberal interpretation of M-Parity, the jokester counts as part of the relevant cognitive process. Recall that the notebook counts as part of the cognitive process by M-Parity because consulting the notebook is an extracranial analogue of biological memory. It's not clear where the jokester fits in<sup>30</sup>. Any plausible version of M-Parity must guard against extending cognition too far, so that the lines between cognitive and non-cognitive processes are erased<sup>31</sup>.

<sup>30</sup> As Clark and Chalmers (1998), p. 8 note, cases of extended cognition occur when the human organism is linked with an external entity in a two-way interaction that creates a *coupled system*. While it's apparent how coupling takes place with the notebook, it's not clear how it could be argued to apply *mutatis mutandis* to the jokester.

<sup>31</sup> See here Spaulding (2012) for a recent presentation of this worry.

#### 4.2 The tension with E-Parity (revisited)

Perhaps the tension between the no-knowledge verdict in Case 1 and M-Parity can be reconciled. Even if it could be, the problem of reconciling the no-knowledge verdict in Case 1 with E-Parity remains. How might we get around the tension here, without giving up the no-knowledge verdict in Case 1 *or* E-parity? The tension, recall, arises in part because it was argued that Case 2 was an intracranial analogue case of Case 1, and that knowledge was present in Case 2. We have already considered in §3.2 why Case 2 is a case of knowledge. Perhaps the place to press, then, is the matter of whether Case 2 really is such that E-parity requires us to count Case 1 as a case of knowledge if we count Case 2 as a case of knowledge. Put another way, perhaps Case 2 is not *really* an intracranial analogue of Case 1. Let's pursue this line.

Consulting a biological memory is indeed the intracranial analogue of consulting the notebook, in Case 1. Though, Case 2 is only the intracranial analogue of Case 1 if we keep the other epistemically relevant features of the case fixed. It was argued in §3.2 that when we did, it turned out the difference between the intracranial and extracranial processes was an epistemic difference, in that we ended up with knowledge in one case, but not the other. And this is just the kind of epistemic bioprejudice that is supposed to be ruled out by E-parity.

Now, one might attempt to argue, for luck-based reasons, that Case 4 (a case where, by the lights of mainstream epistemology, knowledge *is not* present) is more plausibly an intracranial analogue to Case 1 than Case 2 was:

**Case 4:** Otto\*\* tries to remember when his doctor's appointment is today. He comes to believe it is at noon because he remembers an appointment at noon written on a note on the kitchen table. The note on the kitchen table was in fact a note stating that he had a lunch appointment, the following day, at noon. Though his doctor's appointment really is at noon today.

In Case 4, knowledge isn't present. After all, it's just dumb luck that the time of the doctor's appointment was the same as the time of the other appointment written on the note Otto\*\* recalls having seen. Both Case 4 *and* Case 2 can't be intracranial analogue cases to Case 1. After all, knowledge is present in Case 2 but not in Case 4. To emphasise, if Case 4, rather than Case 2, is an intracranial analogue case to Case 1, then the tension with E-parity is avoided. Consider, though, that Case 4 is of the same structural type as Case 5, in that luck undermines knowledge in exactly the same way in Case 4 as in Case 5:

**Case 5:** Shep sees a sheep-shaped object in a field and comes to believe there is a sheep in the field. Unfortunately, what Shep actually saw was a sheep-shaped rock, which he mistook for a sheep. There was, however, an actual sheep hiding behind the sheep-shaped rock, so Shep's belief is true<sup>32</sup>.

<sup>32</sup> This is Pritchard's pet example, which he uses in numerous papers to illustrate Gettier-style (intervening) epistemic luck. Thanks to an anonymous reviewer for pointing out that this case owes originally to Chisholm (1966/1977/1989), not Pritchard.

Recall again that veritic luck comes in *two* varieties, intervening and environmental. That the kind of luck at play in Case 1 was akin to the kind of luck at play in the barn facade case, which is purely *environmental*. What explains why the agent could have easily gotten it wrong in cases of environmental luck is just that the agent is in an epistemically inhospitable environment. Nothing actually goes awry in the belief forming process, as it does in cases of *intervening* epistemic luck.

Now Case 5 is a paradigmatic case where the kind of knowledge-undermining epistemic luck is *intervening* (as was the case in Gettier's (1963) original examples). Given that the structural similarity between Case 1 and the barn facade (environmental luck) case is hard to deny, the prospects for presenting Case 4 (rather than Case 2) as a kind of intracranial analogue to Case 1 are not good. After all, Case 1 is a case where there was no *intervening* epistemic luck, and this structural difference<sup>33</sup> from Case 4 disqualifies it as an intracranial analogue.

## 5 Concluding remarks

I have argued here that proponents of extended cognition should tolerate neither metaphysical *nor* epistemic bioprejudice, and further, that the standard line of thinking about environmental luck in mainstream epistemology seems to betray bioprejudice of both sorts. More specifically, it looks like the no-knowledge verdict in cases where an agent's correctness is down to environmental luck is at odds with both Clark and Chalmers' original parity principle (what I have called metaphysical parity) as well as with the E-parity principle I have argued to be a natural extension of the parity principle in epistemology. This result is hard to accept, and in fact, I don't think we should accept it, at least, not just yet. Because cognitive extension is a relatively new position in the philosophy of mind and the ramifications of this position have not been comprehensively explored in epistemology, the conclusion that either mainstream anti-luck epistemology or extended cognition is just plain wrong, is premature.

I want to conclude by suggesting that the best way to avoid the impasse might be for both anti-luck epistemology and extended cognition to do some re-examining, specifically, of some of their commitments responsible for the impasse I have sketched here. Though anti-luck epistemology seems to get the right result across a spectrum of cases, we need a more precise account of what to hold fixed under the description of the *relevant way the belief was formed* in the actual world, when moving out to nearby worlds. Pritchard (2007) himself has described the account on offer as vague on this point. Anti-luck epistemologists need to do better, and when they do, perhaps this will help deal with cases of extended cognition—cases for which the matter of what precisely to hold fixed is of special importance.

<sup>33</sup> This difference is especially important in Pritchard's recent work, according to which he (e.g. 2012) has argued that intervening epistemic luck, but not environmental epistemic luck, is incompatible with cognitive achievement, where cognitive achievement is understood as a cognitive success that is primarily creditable to the exercise of an agent's cognitive ability. This point has important ramifications for proponents of robust virtue epistemology (e.g. Greco 2010, Sosa 2007, 2009), who take knowledge to be, essentially, a cognitive achievement).

What should proponents of cognitive extension re-examine? Chalmers, in his foreword to Clark's (2011) most recent monograph, *Supersizing the Mind*, reflected on the impact of their co-authored paper which essentially created extended cognition as a research area in 1998. Given the radical nature of the extended mind view, there has unsurprisingly been a great deal of criticism<sup>34</sup> over the past decade or so. As Chalmers sees it, much of the criticism can be effectively dealt with. Chalmers still believes in the extended mind thesis, however, he recognises that the position has a weak spot, which is its collision with plausible insights about the dual boundaries of perception and action. Here's Chalmers:

It is natural to hold that perception is the interface where the world affects the mind, and that action is the interface where the mind affects the world. If so, it is tempting to hold that what precedes perception and what follows action is not truly mental. And one might use this to draw a principled distinction between the cases of Otto (the Alzheimer's patient who uses a notebook as memory) and Inga (the ordinary subject who uses her brain). To interact with his notebook Otto must read it and write in it, requiring perception and action, where there is no such requirement for Inga. If so, then the boundaries above would place the notebook outside the mind. (Chalmers 2011: xi).

The position that environmental luck can undermine the safety of a cognitive success, and disqualify that success as knowledge, relies on some meaningful distinction<sup>35</sup> between the agent's cognitive activity and the environment she finds herself in. When Otto writes in his notebook, he is at once engaging in a cognitive process and, at the same time—through action—extending himself into his environment. Proponents of extended cognition must find a way of extending cognition into the environment while saving the differences.

**Acknowledgments** I would like to express appreciation to Emma C. Gordon, Joel Katzav, Christoph Kelp, Wybo Houkes, Benjamin Jarvis, Philip Nickel, Martin Peterson, Auke Pols, Duncan Pritchard, Krist Vaesen and two anonymous reviewers at Synthese.

## References

- Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology*, *14*(1), 43–64.
- Adams, F., & Aizawa, K. (2008). *The bounds of cognition*. Malden, MA: Blackwell Publishing.
- Carter, J. A. (2010). Anti-Luck Epistemology and Safetys (Recent) Dis- contents. *Philosophia*, *38*(3), 517–532.
- Carter, J. A. (2011). Radical Skepticism, Closure, and Robust Knowledge. *Journal of Philosophical Research*, *36*, 115–133.
- Carter, J.A., Jarvis, B. & Rubin, K. (2012a). Knowledge: Value on the Cheap. *Australasian Journal of Philosophy*. DOI:10.1080/00048402.2012.694455.
- Carter, J.A., Jarvis, B. & Rubin, K. (2012b). Knowledge and the Value of Cognitive Ability. *Synthese*. doi: 10.1007/s11229-012-0220-x.
- Chisholm, R. (1966/1977/1989). *Theory of Knowledge*. Englewood Cliffs, NJ: Prentice Hall.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, *58*(1), 7–19.

<sup>34</sup> See, for instance, Adams and Aizawa (2001, 2008) and (Rupert 2004, 2010).

<sup>35</sup> For a recent survey of arguments on this point, see Spaulding (2012).

- Clark, A. (2011). *Supersizing the mind*. Oxford: Oxford University Press.
- Colao, D., Buckwalter, W., & Stich, S. (2012). Epistemic Intuitions in Fake-Barn Thought Experiments (September 23, 2012). Available at SSRN: <http://ssrn.com/abstract=1973351> or <http://dx.doi.org/10.2139/ssrn.1973351> Accessed 26 Dec 2012.
- Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23(6), 121–123.
- Greco, J. (2010). *Achieving knowledge*. Cambridge: Cambridge University Press.
- Hetherington, S. (forthcoming). There can be lucky knowledge. In M. Steup & J. Turri (Eds.), *Contemporary Debates in Epistemology* (2nd ed.). Malden, MA: Wiley–Blackwell, forthcoming.
- Hetherington, S. (1998). Actually knowing. *The Philosophical Quarterly*, 48(193), 453–469.
- Hetherington, S. (2011). *How to know: A practicalist conception of knowledge*. Oxford: Wiley–Blackwell.
- Hetherington, S. (2012). The extended knower. *Philosophical Explorations*, 15(2), 207–218.
- Jarvis, B. (2012). Knowledge, cognitive achievement and environmental luck. *Pacific Philosophical Quarterly* (forthcoming).
- Kelp, C. (2012a). Knowledge: The safe-apt view. *Australasian Journal of Philosophy*. doi:10.1080/00048402.2012.673726.
- Kelp, C. (2012b). Knowledge, understanding, and virtue. In A. Fairweather (Ed.), *Virtue scientia: Virtue epistemology and philosophy of science*. *Synthese Library*. Dordrecht: Springer.
- Knobe, J. & Nichols, S. (2008). In J. Knobe & S. Nichols (Eds.), *Experimental philosophy*. Oxford: Oxford University Press.
- Kvanvig, J. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge: Cambridge University Press.
- Kvanvig, J. (2004). Nozickian epistemology and the value of knowledge. *Philosophical Issues*, 14(1), 201–218.
- Kvanvig, J. (2008). Critical notice of Pritchards Epistemic Luck. *Philosophy and Phenomenological Research*, 77, 272–281.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Luper, S. (2010). The epistemic closure principle. *The Stanford Encyclopedia of Philosophy* (Fall 2010 Edition). <http://plato.stanford.edu/archives/fall2010/entries/closureepistemic/>. Accessed 26 Dec 2012.
- Madison, B. J. C. (2011). Combating anti anti-luck epistemology. *Australasian Journal of Philosophy*, 89(1), 47–58.
- McDowell, J. (1994). *Mind and world*. Harvard: Harvard University Press.
- Peterson, M. (2009). *An introduction to decision theory*. Cambridge: Cambridge University Press.
- Plantinga, A. (1993). *Warrant and proper function*. Oxford: Oxford University Press.
- Pritchard, D. (forthcoming). There cannot be lucky knowledge. In M. Steup & J. Turri (Eds.), *Contemporary debates in epistemology* (2nd ed.). Malden, MA: Wiley–Blackwell, forthcoming.
- Pritchard, D. (2003). Virtue epistemology and epistemic luck. *Metaphilosophy*, 34, 106–130.
- Pritchard, D. (2004). Epistemic luck. *Journal of Philosophical Research*, 29, 193–222.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: OUP.
- Pritchard, D. (2007). Anti-luck epistemology. *Synthese*, 158, 277–297.
- Pritchard, D. (2010). Cognitive ability and the extended cognition thesis. *Synthese*, 175, 133–151.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109, 247–279.
- Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *The Journal of Philosophy*, 101, 1–40.
- Rupert, R. (2010). *Cognitive systems and the extended mind*. Oxford: Oxford University Press.
- Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. 1). Oxford: Oxford University Press.
- Sosa, E. (2009). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. 2). Oxford: Oxford University Press.
- Spaulding, S. (2012). Overextending cognition. *Philosophical Psychology*, 25(4), 469–490.
- Steup, M. (2008). The analysis of knowledge. *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition). <http://plato.stanford.edu/archives/fall2008/entries/knowledgeanalysis/>. Accessed 26 Dec 2012.
- Swain, M. (1978). Reasons, causes, and knowledge. *Journal of Philosophy*, 75, 229–249.
- Zagzebski, L. (1996). *Virtues of the mind*. Cambridge: Cambridge University Press.