

Extended Cognition and Propositional Memory

J. Adam Carter & Jesper Kallestrup

University of Edinburgh

Abstract. The philosophical case for extended cognition is often made with reference to ‘extended-memory cases’ (e.g. Clark & Chalmers 1998); though, unfortunately, proponents of the hypothesis of extended cognition (HEC) as well as their adversaries have failed to appreciate the kinds of epistemological problems extended-memory cases pose for mainstream thinking in the epistemology of memory. It is time to give these problems a closer look. Our plan is as follows: in §1, we argue that an epistemological theory remains compatible with HEC only if its epistemic assessments do not violate what we call ‘the epistemic parity principle’. In §2, we show how the constraint of respecting the epistemic parity principle stands in what appears to be a *prima facie* intractable tension with mainstream thinking about cases of propositional memory. We then outline and evaluate in §3 several lines of response.

Keywords: Extended cognition; propositional memory; knowledge; cognitive ability

§1 Extended Cognition and Epistemic Parity

Consider the following case, due to Clark & Chalmers (1998):

Inga: Inga relies on her biological memory in a way that is perfectly normal; new information she learns is stored in her biological memory, which is also what she consults when she needs some old information.

Two obvious observations are in order. First, Inga’s biological memory is part of the cognitive process she employs in the ordinary course of storing and retrieving information. Secondly, Inga’s biological memory is *intracranial* in that it is seated within her skin-and-skull.

But now imagine that Inga’s biological memory were instead somewhere *outside* her skull and skin, e.g., enclosed in a brain-in-a-vat that she carries around with her, playing the same role *vis-à-vis* stored and retrieved information as it usually does, intracranially, in the default case. While this would be a strange sight, we would surely not hesitate to include her portable biological brain as part of the cognitive process she employs in remembering the kinds of things she remembered before; her biological memory is after all playing the very same functional

role—*viz.*, the role of being that which stores the new information she learns and the old information she retrieves. It would just be doing so in a way that goes beyond her bodily boundary. So, mere location can't be what matters. What then about implementation? Imagine that part of Inga's biological memory, say her dorsal premotor cortex, is replaced in the same location with a silicon-chip device which plays the same functional role in terms of storing and retrieving information. Again, we would not be reluctant to include that artificial implant as part of the cognitive process that she uses in remembering events that occurred before the surgery. So, actual constitution can't be what matters either. But now consider a combined case where part of a cognitive process is both constituted non-biologically and occurs outwith the skin-and-skull. Here is Clark & Chalmers's (1998) case of 'Otto':

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.

Is Otto's notebook part of the cognitive process he employs? Despite the strong intuitive pull to think that Inga's biological brain is part of her cognitive process—and would remain so even when extruded from her skull into a portable vat, or when replaced with a silicon chip—a common knee-jerk reaction is to insist that Otto's notebook is *not* a part of the cognitive process he employs.

According to proponents of the *hypothesis of extended cognition* (HEC)¹ this knee-jerk reaction to disqualify Otto's notebook as part of the cognitive process he employs reveals an unwarranted 'metaphysical' *bioprejudice*²—an unprincipled bias that features in our judgments about the metaphysical nature of a cognitive process. As Clark (2011) notes, such prejudice is exercised when it is precisely *because* something is non-biological and extracranial that it is excluded as part of a cognitive process. Put differently, *were it biological and intracranial we would not hesitate to include it as part of a cognitive process.*

¹ HEC is one of several varieties of what has been called 'active externalism'. Unlike 'passive', content externalism (*a la* Putnam (1975) and Burge (1986)), active, or vehicle, externalism allows for proximate elements in the world to feature in cognition. For a sample of proponents of various forms of active externalism, see (Clark & Chalmers (1998); Clark (2007; 2008); Hutchins (1995); Menary (2006; 2007); Theiner (2011); Wheeler (2010), Wilson (2000; 2004).

² This is Clark's (2011, *xviii*) term.

Originally presented in Clark & Chalmers (1998), the much-referenced *parity principle* is designed to guard against such (metaphysical) bioprejudice:

(METAPHYSICAL) 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 & Chalmers 1998, 8).

Because Otto's notebook plausibly plays the kind of functional role *vis-à-vis* information storage and retrieval as Inga's biological memory, then, by reference to the *parity principle*, because Inga's biological memory is ruled-in as part of the cognitive process *she* employs, so Otto's notebook is ruled-in as part of the cognitive process *he* employs. This is at any rate the position advanced by Clark & Chalmers, which has since become a defining idea used to motivate the extension of cognitive processes.³

But for now, we will shift tack in a way that brings the relevance of the position sketched thus far into focus *vis-à-vis* issues in mainstream epistemology. Epistemological theories are in the business of making *epistemic assessments*. Clearly, such theories can commit a kind of markedly 'epistemic' bioprejudice in the course of making epistemic (rather than metaphysical) assessments. For example, a theory of knowledge *T* reveals a kind of bioprejudice (at odds with HEC) if, by the lights of *T*'s own epistemological criteria, a case that *otherwise* would be ruled-in as knowledge, is ruled-out just on account of the fact that part of the process employed is non-biological.

A commitment to respecting Clark & Chalmers' original (metaphysical) parity principle will not go very far toward safeguarding against this kind of 'epistemic' bioprejudice. A new kind of principle is needed: an *epistemic parity principle*⁴.

EPISTEMIC PARITY PRINCIPLE: For individual *S* and belief in *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*.

By reference to the epistemic parity principle, we're in a good position to investigate what kinds of epistemic assessments stand in tension with HEC.

³ Of course, one may challenge the claim that Otto's notebook and Inga's biological memory play identical functional roles in the relevant sense. We address this concern at the end of §3.1.

⁴ See Carter (2013) for a more detailed discussion of this principle; *Cf.* Carter & Pritchard (2014).

§2 Epistemic Parity and Propositional Memory

The requirement to uphold the epistemic parity principle generates new puzzles with respect to the epistemic assessments made in cases of *propositional memory*. By ‘propositional memory’, which will be our focus in what follows, we are referring to the variety of memory naturally given expression in the form of ‘S remembers *that* p,’ where p is a proposition.⁵ Ultimately, as we’ll see, respecting the epistemic parity principle turns out to be more difficult the more we think about Otto’s case from the lens of mainstream theorizing about propositional memory.

§2.1 Memory Causation and External Loops

Consider now, following Bernecker (2010), a kind of ‘template’ outlining some basic necessary conditions for a subject to count as *remembering* a given proposition, such as that MoMA is on 53rd street:

Propositional memory: S remembers at t_2 that p only if:

- (1) S represents at t_2 that p ,
- (2) S represented at t_1 that p^* ,
- (3) p is true at t_2 ,
- (4) p is identical with, or sufficiently similar to, p^* ,
- (5) And S’s representation at t_2 that p is suitably causally connected to S’s representation at t_1 that p^* .

Some quick points of caution. These conditions are not meant to constitute a substantive account of propositional memory; different substantive accounts will be functions of how (1)-(5) are filled out. Secondly, ‘representation’ is used in this template in an inclusive way so as to be compatible with, though not presuppose, the *epistemic theory of memory*⁶. According to this theory, the representational relations that feature in (1), (2) and (5) are, necessarily, *knowledge (or justification) relations*⁷. Granted, in the default case, the representations at issue will often meet

⁵ For simplicity’s sake, our focus will be cases of *extroversive* propositional memory, setting aside *introversive* cases. See here Bernecker (2012: 21).

⁶ For some notable defenses of the strong epistemic theory of memory, see Ryle (1949), Ayer (1956), Pollock (1974) and Williamson (2000) and, more recently, Moon (2013).

⁷ See here Norman Malcolm’s suggestion that one remembers a proposition just in case he knows it *because* he knew it. (Cf. Bernecker 2012: 66). The epistemic condition plays a causal role: not only must the representations that feature in (1), (2) and (5) constitute knowledge, one’s previous knowledge must also cause one’s present knowledge.

knowledge conditions. But the epistemic theory of memory is making a stronger claim meant to apply to all cases of *bona fide* memory. This view has been a controversial one, even on a weaker ‘justification’ interpretation,⁸ and we’ve no reason for our purposes to simply assume it *ex ante*. Furthermore, to appreciate some of the new philosophical problems extended-cognition cases pose for mainstream thinking about propositional-memory, it will be more instructive to leave open the matter of how to articulate the relevant representations in cases of *bona fide* memory.

These qualifications made, let’s revisit (1)-(5), bearing in mind that these conditions are meant to be minimal necessary requirements on any viable account of the epistemology of memory. As Bernecker notes, the first two conditions are supposed to uphold the idea that memory involves retention; condition (3) preserves factivity in that remembering that *p* entails the truth of *p*. Condition (4) ensures that one counts as remembering only something that was previously represented, and importantly for our purposes, condition (5) precludes *relearning*.

A blatant example of relearning, which is incompatible with remembering, occurs when Tim represents *p* at t_0 , forgets *p* at t_1 , and then comes to represent *p* anew at t_2 . In such an instance, even though Tim represents a content he had previously represented, there clearly isn’t any causal connection between his representing *p* at t_0 and later again at t_2 —and so (5) won’t be met.

A live philosophical question is how to get the relevant causal connection right. The details for how to make the causal condition *sufficient* are debated amongst proponents of the causal theory of memory⁹. It’s beyond our present scope to adjudicate between these proposals. What is pressing here is that Otto must at *least* satisfy some plausible construal of (5)—or else Otto’s case is not a case of propositional memory, extended or otherwise.

On a first pass, one might opt for a simple *counterfactual approach* to the causal condition in (5)¹⁰: for instance, with reference to our case of Tim, the idea would be that Tim’s current representation at t_2 that *p* is suitably causally connected to his originally representing *p* at t_0 only if his representing *p* at t_2 counterfactually depends on his representing *p* at t_0 such that: had he not represented *p* at t_0 , he would not represent *p* at t_2 .

While this simple view can handle easy cases, it runs into notorious problems¹¹ with cases that, like Otto’s, involve *external loops*—*viz.*, where part of the relevant causal chain occurs *outside* the biological individual. A famous example of such a case is originally due to Martin & Deutscher (1966):

⁸ See here Bernecker (2012: Ch. 3) and Matthen (2010).

⁹ See here Bernecker (2008; 2010). Martin & Deutscher (1966).

¹⁰ See, for instance, Lewis (1973). *Cf.* Paul and Hall (2013).

¹¹ For a classic example of such a case, see Martin & Deutscher (1966: 180-181).

Suppose that at t_1 S witnesses a car accident. At t_2 , S tells his friend all about the accident. Then, at t_3 , S gets a blow on the head as a result of which he permanently forgets everything about the accident he witnessed at t_1 . When, at t_4 , the friend notices that S can no longer remember the car accident, he tells S the details that S had told him at t_2 . S hears the friend tell him about the car accident and, at t_5 , this causes S to seem to recall the accident via a causal chain that includes his friend's memories and honest intentions¹².

Now *even* though it's evident that a causal chain supports a counterfactual correlation between what S observes of the accident at t_1 and S's representing the accident at t_5 , it's utterly implausible to say that S *remembers* the accident¹³. Obviously, one might be tempted to say that the problem with the causal chain in the car-wreck case is simply *the external loop* element—an observation that suggests we need merely supplement the simple counterfactual requirement with the stipulation that the entire causal chain take place *within* the biological individual.

But merely stipulating that external loops are incompatible with memory will *in principle* exclude Otto as a candidate for memory *ex ante*—after all, part of the causal chain between Otto's storage and recall is external to his biological body. Dealing with the car-wreck case by simply disbaring external causal loops as candidates for satisfying (5) thus rules out from the outset the very possibility of extended memory¹⁴.

Importantly, though, the reason to not exclude external loops *as such* is also independently motivated by considerations that have nothing to do with extended cognition, *per se*. The worry is that such exclusion would be an *ad hoc* and overly strict constraint on memory. Why? The rationale here is just a simple insight of multiple realizability. Here is Bernecker (*Ibid.* 129-130):

It is surely a contingent fact that remembering depends on processes inside the nervous system. It is conceivable that information is stored in parts of the brain that can be detached, stored and reattached to the brain... the spatial location of the causal chain is inessential.

¹² This is Bernecker's (2010: 129) concise paraphrase of the original case.

¹³ See Bernecker (2010: 129-130).

¹⁴ Bernecker (2010: §6) is not viewing cases such as 'detachable external storage' as examples of either extended mind or extended cognition.

Since simply prohibiting external loops, *tout court*, is unprincipled, then—some further causal criterion must be adverted to in order to (i) rule out obvious cases of relearning, such as we find in the ‘car wreck’ case; while (ii) not in principle ruling out what (by the lights of HEC) will be *benign* external loop cases, such as Otto’s.

Bernecker suggests the way forward is, put simply, to bring on board a kind of contiguous trace condition, according to which S’s representation at t_1 that p^* and S’s representation at t_2 that p are connected by (i) a *persisting memory trace* or (ii) a *contiguous series of memory traces*¹⁵—where traces, in their role of preserving content are dispositional (or non-occurrent) beliefs in the sense of information available to the mind for endorsement¹⁶. Obviously, in the car-wreck case, there is no persistent memory trace, nor a contiguous series of memory traces (e.g. dispositional beliefs). For Otto, it’s not at all obvious that there is *not* a persistent dispositional belief¹⁷, so this suggestion has some initial promise. We’ll return to this point in due course.

§2.2 Extended Memory—A Closer Look

Before engaging with the issue of whether Otto can be said to have a contiguous series of memory traces by examining whether he counts as having the relevant dispositional beliefs, it’s important to address first what appears a more flagrant problem with Otto as a prospective candidate for extended memory—namely, it looks like Otto is relying on an aid or prompt in a way that is in *prima facie* tension with any plausible satisfaction of the contiguous trace condition.

¹⁵Bernecker (2010: 130). Memory traces are understood to have a dual role: preservation of mental content through time and production of states of recall (*Ibid.* 131). As regards the latter role, memory traces are typically given a physical description (*Ibid.* 132). The mental description—in terms of belief states—applies to the former role. We will be focusing on the mental description here, in part because this is of primary philosophical interest, and accordingly will be most illuminating for seeing how Otto might satisfy (5); furthermore, we will later be parting ways with Bernecker’s intracranial account of memory trace creation and memory state production as ‘inside the head’ (*Ibid.* 181).

¹⁶To dispositionally (or non-occurrently) believe is different from having a disposition to believe. When S dispositionally believes p , S can access p by a retrieval process from memory, but when S has a disposition to believe p , p isn’t preserved in memory at all, and so S can access p only by forming an occurrent belief that p . For more on this distinction see Audi (1994). Note that the case of Otto was designed by Clark & Chalmers to be guided by intuitions that derive from common-sense talk of dispositional belief. Note also that we are still focusing on propositional memory. In the case of non-propositional memory, memory traces are sub-doxastic states (*Ibid.* 130).

¹⁷Or, at least, a contiguous series of dispositional beliefs, the original being intracranial, and the intracranial causing the extracranial (e.g. Otto’s occurrent belief causes him to write in the notebook what, by HEC, constitutes a kind of extended dispositional belief.)

To get a feel for this worry, just consider the following objectionable example of relearning, where a prompt/aid is simply parroted¹⁸:

Memo: Memo once truly represented the names of all the students in his primary school class; but because he decided it would be too taxing to keep these names in biological storage, he bought a yearbook. Years later, Memo accidentally opens the yearbook and represents p (a proposition consisting in the names of the students).

Memo obviously doesn't remember the names of his primary school students, and the flagrant foul here seems to be that Memo is relying on an aid *entirely*. But here we must be careful. *What* about Memo's reliance on the aid is objectionable? For a first pass it's natural to think that what's objectionable is just that the prompted information is identical to the stored information; unlike, for instance, in cases where you have your memory triggered by a cue, after which you reproduce *more* information than was supplied by the prompt.¹⁹ Contrast, though, with this case:

Joggo: Joggo normally truly represents his computer password, but now he suddenly cannot reproduce the four digits. When the IS Helpdesk informs him that his password is 1245, he immediately recognizes it: 'Oh yeah, that's right 1245'.

Here Joggo exhibits what Martin & Deutscher, as well as Bernecker, call 'complete prompting'—in that the prompted information and the stored information are the same. But even so, the Joggo case is indeed very different from the Memo case—there is a clear sense in which Joggo's memory is triggered by a cue here *despite* the fact that the prompted information is identical with the stored information. Joggo remembers, but Memo doesn't. Plausibly, a relevant difference for why we think Joggo's memory is triggered by a cue, but Memo's is not, is that Joggo (unlike Memo) is such that the prompt actually *causes* Joggo's current representation to be

¹⁸ Although Bernecker, along with Martin & Deutscher, allow what they call 'complete' or 'strict' prompting' to be compatible with genuine memory, Bernecker equates 'merely parroting the prompt' with relearning.

¹⁹ Since we take the content of a representation as well as information itself to be a proposition p , which is an abstract entity, we shall take the *causal relata* to be the concrete vehicles that run propositional content, e.g. S's particular representation that p or an informational token state with p as its content.

suitably causally connected to his original representation²⁰. Put in terms of the contiguous trace constraint: the prompt plays the causal role for Joggo of causing Joggo's current representation to be caused by his original representation, which for him will have been a dispositional belief that Memo seems to lack. Let's call the unobjectionable kind of prompting cases—like that of Joggo—'triggering cue cases'²¹:

Triggering Cue: (i*) the prompted information need not be (but can be) identical with the content of S's original representation, and (ii*) the state of the prompted information causes S's current representation to be suitably causally connected to S's original representation.

Obviously, if Otto is in the same boat as Joggo, there's no principled barrier to retaining the suggestion that Otto counts as remembering—as triggering cue cases will (thanks to (ii)) de facto *satisfy* the 'trace condition', and in doing so, plausibly constitute cases that feature a suitable causal connection between the originally stored content and the currently represented content (as required in (5)).

Problem is, though, there are some clear differences between Otto and Joggo that make it difficult to defend Otto's case as one of a triggering cue. One apparent difference between Otto and Joggo is that, in the case of Joggo, it seems intuitively easy to see how the state of the prompted information causes Joggo's current representation to be suitably causally connected to his original representation (recall: this ensures condition (ii) of *triggering cue*). Things are, at least *prima facie*, quite a bit less straightforward for Otto. One way to capture this point is to appeal to a certain phenomenon in the causation literature called *double prevention*. As Peter Menzies summarises the phenomenon:

Double prevention: A person is walking along a mountain trail, when a boulder high above is dislodged and comes careering down the mountain slopes. The walker notices the boulder and ducks at the appropriate time. The careering boulder causes the walker to duck and this, in turn, causes his continued stride. (This second causal link involves double prevention: the duck prevents the collision between walker and boulder which, had it occurred, would have prevented the walker's continued stride.) However, the

²⁰ Thanks to Duncan Pritchard for suggesting that a further difference here is that Joggo's ability to recall his password in the days or weeks following prompting is more robust, other things being equal, than is Memo's ability to recall his classmates' names.

²¹ This condition is meant to be general enough to encompass both strict and complete prompting cases.

careering boulder is the sort of thing that would prevent the walker's continued stride and so it seems counterintuitive to say that it causes the stride.²² (2014: §3.3)

By parity of reasoning, it looks like Otto's original, and seemingly defective, memory effectively *prevents* his having a persisting dispositional belief, which, *had he had one*, would have prevented him from relying on the notebook to trigger his memory²³. However, a defectively stored memory is ordinarily the sort of thing that would prevent one from having one's memory triggered, and so it seems counterintuitive to say that Otto's original encoding of the memory causes the triggering. To the extent that double prevention cases are effective in eliciting the intuition that Hall (2000) originally thought they did, we might thus find ourselves having similar intuitions about Otto that would lead us to deny that he satisfies condition (ii) of *triggering cue*²⁴.

The issue of double prevention aside (we'll return to this), there is also a phenomenological point to make; it's not initially clear how to make sense of 'triggering' in the absence of what Russell (1921: 186) calls a certain 'specific feeling... [or] complex of sensations' that would accompany a triggered memory. Otto does not express 'oh yeah, that's right' when consulting his notebook.²⁵ The *phenomenology* of the case of Otto is much more like the case of Memo than that of Joggo²⁶.

²² This example owes originally to Hall (2000)'s criticism of the transitivity element of Lewis's causal chains.

²³ At least, Otto's belief (unlike Joggo's) is *prima facie* defectively stored, in the sense that the original storage location (biological memory) for Otto is failing and would not by itself causally sustain the information.

²⁴ It might be thought that Joggo fails to satisfy condition (ii) of triggering cue no less than Otto. However, this worry is premature. This point can be made in two ways. Firstly, the double prevention worry relies on the intuition that the original storage is defective, which is clearly in play in the case of Otto given that his original biological storage source would not sustain the information. By contrast, Joggo's original storage source does sustain the information in such a way that it can be suitably causally connected, *via* that storage source, to Joggo's current representation when his memory is triggered. Secondly, if *ex hypothesi* Otto's memory were triggered, it would be triggered in a manner very differently from Joggo's (and in a way that would be relevant vis-à-vis the double preventing intuition). Consider that, as Bernecker (2010: 140-141) notes, memory states are often joint products of a trace and an appropriate prompt or retrieval cue, where the latter can be either controlled or automatic. While the apparent defect in Otto's original (biological) memory storage seems to lead him to engage in the controlled process of making use of the notebook, the same can't be said for Joggo, whose triggering is rather of the kind that is automatic, where the cue directly facilitates the retrieval of the target representation.

²⁵ See here also Senor (2009).

²⁶ Of course, an advocate of HEC could reply that the phenomenon of triggering mustn't be specified in such a way as to rule out the possibility of extended-triggering, *ex ante*, by making a biologically-grounded kind of phenomenology *essential* to triggering (e.g. by aligning triggering with a kind of triggering quale).

§2.3 A Problem for Preserving Parity

The problem of countenancing Otto's case as a case of extended *memory* only magnifies once we revisit the *epistemic parity principle*, as outlined in §1. Recall that, by the lights of this principle, Otto counts as having memory knowledge (i.e. knowledge supported by memory) so long as Inga (who relies entirely on biological memory) does, where all epistemically relevant details are held fixed, except that Otto is using the extended analogue of the process Inga uses. To bring the problem here into sharp relief, let's focus on a single proposition: *that MoMA is on 53rd street*. Now compare again the Otto and Inga cases from §1. Suppose that each recalls that proposition, Inga by consulting her biological brain, where this information is stored for her, and Otto by consulting his notebook, where it is stored for him.

Now here's the crux: either Otto has memory knowledge, or he does not. One can't very well preserve the epistemic parity principle while at the same time denying that Otto counts as attaining memory knowledge, when Inga does. And so insofar as epistemologists would wish to avoid ruling out HEC, it's incumbent on them to vindicate how Otto has such knowledge no less than Inga does.

But there are two rather deep problems here. The first is that, obviously, Otto lacks memory knowledge if what he exhibits is not what Inga exhibits. As we saw in §2.1, how Otto can be vindicated as having genuine memory is yet to be shown, given the superficial causal similarities with Memo. The second problem is that *even if* it can be explained how Otto counts as remembering despite seeming to merely parrot the prompt as Memo does, there is an *independent reason for thinking Otto lacks knowledge*. The argument here is that Otto seems to fail to satisfy a kind of platitude in epistemology, which is that knowledge requires some suitable exercise of cognitive ability²⁷.

²⁷ For a number of notable explicit defenders, see Sosa (1991; 2007; 2009), Greco (1999; 2010), Pritchard (2010; 2012; 2013.) The tacit support for this insight runs much wider; both internalists and externalists about epistemic justification, for instance, will be inclined to deny knowledge in cases in which cognitive ability is not connected with the production of a true belief. One way to appreciate this point is *via* a simple thought experiment. Suppose, following Pritchard (2010: 222), that Sue is—entirely unbeknownst to her—being 'fed' beliefs about her environment by supercomputers. Sue might very well form true beliefs, which are also modally stable across a range of possible worlds (and thus they could be both safe and sensitive); these beliefs might also be backed by good reasons, and be epistemically blameless. But these beliefs would surely not count as knowledge despite their epistemic credentials, and the reason is that their truth has nothing to do with Sue exercising her cognitive abilities. Likewise, it is a lack of exercise of ability that is a common explanation for why Lehrer's (1990) case of 'TrueTemp'—a famous counterexample to process reliabilism—is not a case of knowledge. Lehrer's case features an agent who compulsively (thanks to a special thermometer that is, entirely unbeknownst to him, implanted in his brain) forms true beliefs about the current ambient temperature. Though TrueTemp's beliefs are reliable,

To make this worry more concrete, consider that if the truth of S's belief is either down to some environmental happenstance or through the exercise of someone else's ability, then that belief cannot constitute knowledge. Consider the case of Freya:

Freya: Suppose Freya has no ability to reliably detect ordinary room temperatures. Out of the blue, Freya forms the belief that the temperature in a particular room is 20°C. In one case that just happens to be true. In another case, some hidden jokester unwittingly adjusts the room temperature so as to perfectly match Freya's belief.²⁸ Because in each case the truth of her belief is disconnected from any ability of hers, she lacks knowledge.

Freya's correctness is certainly not in any clear way down to a cognitive ability we can plausibly attribute to *her*. With reference to the platitude that knowledge requires the exercise of cognitive ability on behalf of the individual, Freya therefore fails to know.

In one respect, it's both tempting and correct to say that Memo, despite not remembering, is much better off than Freya. Memo after all has *perceptual* knowledge from looking at the yearbook. But that's not what's at issue. Memo clearly does not have *memory* knowledge. That is, even if we granted that Memo's case was one of genuine memory (an assumption not even the most ardent friends of HEC would allow), it would remain that Memo's correctness is entirely unrelated to any *memorial* ability on his part. Given the superficial similarities between Memo and Otto, the worry then becomes that even if we set aside the previous concerns and grant that Otto counts as attaining *bona fide* memory, one can nonetheless insist that Otto, in light of his reliance on the notebook, is not in the market for *memory* knowledge. His getting it right isn't suitably down to a *memorial* ability we can reasonably credit to him²⁹.

he is near-unanimously regarded as lacking knowledge, and an obvious explanation is that his correctness is unconnected to any cognitive ability we can attribute to *him*. Furthermore, the explanation for why we cannot attribute TrueTemp's cognitive success to *his* cognitive ability is, as Greco (2010) has noted, because the thermometer is not appropriately *integrated* into TrueTemp's cognitive agency.

²⁸ We could even imagine that the jokester's environmental interventions are highly reliable so that not easily would Freya form false temperature-beliefs. Despite the safety of her beliefs, intuition has it that she would still lack knowledge. This twist on the example illustrates that knowledge is governed by two master intuitions to do with ability and safety (or perhaps sensitivity), which can arguably come apart.

²⁹ The ability condition on knowledge is not to be understood so weakly that *any level of ability* counts as satisfying the condition. The idea is rather that the correctness of one's belief must be at least to some significant extent creditable to the exercise of ability. Following Greco (2009: chap. 5), when S attains knowledge the overarching, or most salient, element in a causal explanation of her cognitive success is her cognitive ability. Thus even if Otto exercises some non-negligible ability in complete prompting cases, this

§3 Strategies of Response

Respecting epistemic parity, then, will involve allowing that Otto has memory knowledge—given that we attribute this to Inga—and then engaging with both associated worries: one which challenges the Otto case as one of *bona fide* remembering, the other which challenges whether Otto counts as *knowing* regardless of whether the case can be defended as one of propositional *memory*. We'll now look at what we take to be the most promising strategy for developing such an argument: a strategy that involves, in an important respect, *turning-the-tables*.

In this section, we will explore a bold way out of the jam sketched to this point—a strategy according to which the case of Otto is claimed to resemble in all epistemically relevant respects neither Memo nor Joggo, but *Inga*—except that whereas Inga's knowledge is supported by her biological memory, an extended memorial process supports Otto's knowledge³⁰. We'll now turn to considering the prospects of this proposal in some detail.

§3.1 Extended Memory

It's interesting to note that the worries that Otto fails the relevant causal condition (5) are worries that are, in different ways, premised upon a conception of Otto's memory as, in some relevant respect, exhausted by his (failing) biological memory. The turn-the-tables line will attempt to nip issues premised upon this conception off at the pass.

Firstly, consider that the double prevention worry—which was alluded to in suggesting that Otto's case is not one of triggering—also doubles as a more general worry that there isn't an appropriate causal chain in place. But this worry looks considerably less credible once we take the claim made by proponents of HEC seriously. That is, once (by reference to the parity principle) we accept the view that Otto's notebook is playing the same role for him as Inga's biological memory is playing for her—then the result is that Otto's case is *no more a case of double prevention than is Inga's*. Obviously, Inga's case is *not* one of objectionable double prevention. Therefore, neither is the case of Otto; or so the argument goes.

By the same token, recall in §1 that a worrying issue seemed to be that Otto appeared to *rely on an aid* like Memo, in a way that seemed akin to merely parroting a prompt—which is

won't be enough to satisfy the condition as ordinarily understood. See also Pritchard (2012) and Kallestrup & Pritchard (2012; 2013).

³⁰ In reply to an anticipated phenomenologically grounded objection: note that in the absence of an independent argument to the effect that there is no extended analogue to the biological phenomenology in the default case of memory knowledge, the *mere fact* that the phenomenologies are distinct will be dismissed by advocates of HEC as question-begging.

relearning rather than remembering. It was then suggested that to the extent Otto resembles Memo in merely parroting a prompt, we have grounds to doubt that Otto satisfies what Bernecker (2010, 130) called the *contiguous trace condition*—that S’s representation at t_1 that p^* and S’s representation at t_2 that p are connected by (i) a persisting memory trace or (ii) a contiguous series of memory traces. Why *exactly* does it seem that Otto’s apparent parroting the prompt is in tension with this condition? Perhaps the reasoning is supposed to go like this:

- (1) Otto’s consulting his notebook causes him to represent the relevant content p at t_2 .
- (2) If Otto’s consulting his notebook causes him to represent the relevant content at t_2 , then Otto’s representation that p at t_2 is not connected to his representing p^* at t_1 via a persistent memory trace or a contiguous series of memory traces (*Contiguous trace condition*)
- (3) Therefore, Otto’s representation that p at t_2 is not connected to his representing p^* at t_1 via a persistent memory trace or a contiguous series of memory traces [1,2]

Obviously, (1) is true, but this needn’t be any sort of problem for Otto so long as the notebook is part of a contiguous series of memory traces—a condition Otto satisfies of course if the entry stored in the notebook constitutes a *dispositional belief*. Trickily here, though, if Otto’s entry *is* playing the relevant role of a dispositional belief, then it’s easy to use this point as a premise in a simple deductive argument to the conclusion that Otto satisfies the contiguous trace condition. In short, (1) generates a problem for Otto *only* if one already is committed to denying that Otto will satisfy this condition.

Interestingly, (2) draws one to the same issue: *whether* the notebook’s causing the entry should count against Otto’s satisfaction of the contiguous trace condition turns on one’s previous conception of whether Otto counts as having a dispositional belief—*viz.*, that the notebook entry that causes his representing p at t_2 itself constitutes a dispositional belief.

Bernecker has weighed in explicitly on this point, though in a way that must be placed in context. As we noted earlier, he (2010: 129-130) is happy to allow for the kind of scenario Martin & Deutscher envisioned to qualify as genuine memory—*viz.*, that brain-bound information can be detached into external storage and then reconnected to the brain. But in granting this Bernecker immediately takes care to distance himself from the extended mind thesis (EMT). In doing so (2010: §5.4) he engages directly with the case of Otto. Context is important here because (as is often overlooked) Clark & Chalmers conceived of the Otto case to be motivating via common-sense functionalism *both* HEC and the comparatively more radical EMT. Bernecker

does not discuss extended cognition *per se*, but he does explicitly address the matter of whether Otto might have a kind of dispositional belief. He remarks:

When [Otto is] asked at t_1 whether he believes that the museum is on 53rd street [Otto] cannot say that his notebook believes that the museum is on 53rd street. For if [he] knew his notebook's beliefs all along, then he wouldn't need to consult the notebook. But then why should we attribute to [Otto] the belief that the museum is on 53rd street? The belief attribution seems to be ungrounded.

The turn-the-tables reply, in response to this kind of reasoning, begins with the premise that Otto's notebook is functionally isomorphic to Inga's biological memory, and then replies along the following lines: when Inga is asked at t_1 whether she believes the Museum is on 53rd street, Inga will not say that her biological brain believes the Museum is on 53rd street. In this respect, she is in the same position as Otto in that neither identifies what plays the functional role of storage with what believes the stored information. The turn-the-tables reply continues: Just as it is illicit to suggest that if Inga knew where the Museum was all along she wouldn't have to consult her biological brain, neither is it licit to say that if Otto knew where the Museum was all along he wouldn't have to consult his notebook. Finally, the attribution of a dispositional belief to Otto—which is, again, all that is needed for Otto to satisfy the relevant contiguous trace condition—isn't tantamount to a full-blown endorsement of EMT. All it takes for S to dispositionally believe p is for a representation with p as its content to be stored in S's memory such that S would endorse p when retrieved. Even if Bernecker is right to reject EMT (an issue we will not address), this doesn't constitute an argument against the kinds of dispositional beliefs that are entailed by HEC, and which when in place enable Otto's memorial process to be upheld as on a functional par with Inga's.

Of course one may take issue with this last claim. Thus Rupert (2004: 407-424) has raised the following *dilemma* for the proponent of HEC. If she opts for fine-grained functionalism then HEC is false, because there are fine-grained differences between the functional roles of biological memory and extended memory. For instance, interference and generation effects are widely exhibited in ordinary biological memory, but are absent in the case of Otto. However, if she settles for coarse-grained functionalism involving a generic conception of memory, then there is little if any causal-explanatory work for HEC to do in cognitive science. Instead Rupert (2004; 2009) advances a hypothesis of *embedded* cognition (HEMC) according to which the human cognitive system is organismically bounded, but external

resources causally influence cognitive processing in surprising ways without constituting such processing.

Clark's (2008: 91-92; 2010: 44) response to the first horn is that if cognition is delineated by such fine-grained features then intelligent Martians and other creatures whose physical makeup differs from ours are also barred from memorial processes. Surely, such creatures could have a psychology that differs from ours in some fine-grained manner, e.g. by not being subject to negative transfer, without thereby lacking memory. Clark (2008: 88, 115-121; 2010: 50-51) counters the second horn by arguing that a mature cognitive science is bound to identify a unified causal-explanatory role, as implicitly grasped by normal human beings, for a *broad* cognitive kind of memory to play. For the integrated operations between the internal biological makeup and the external resources are so complex and intricate that no piecemeal decomposition of the extended system will provide a satisfactory understanding of its behavior.

Since delving deeper into either response would take us too far afield from the epistemological questions we primarily seek to address, consider instead Chalmers & Clark's (1998: 10-13) claim that the kind of memory, which their common-sense functionalism pertains to, plays an indispensable role in belief-desire explanations of behavior. Thus, on HEC a simple explanation of why Otto walked to 53rd Street can be provided in terms of Otto wanting to go to MoMA and believing, even before consulting the notebook, it was on 53rd Street. So, HEC allows for Otto's behavior to be subsumed under a familiar folk psychological kind: believing that MoMA is on 53rd street. Denying HEC means any explanation of Otto setting off will be needlessly complex: Otto desires to go to the Museum, he believes the Museum is located at the address in the notebook, and he has access to the fact that the notebook says the Museum is on 53rd Street. On any alternative account such as HEMC, Otto's prior belief that the Museum is at the location written in the notebook leads him to the looking which in turn leads to the formation of an occurrent belief about the actual street address. Clark (2010: 46) dubs this 'the Otto 2-step'. Importantly, the latter belief then disappears as soon as the notebook is filed away. Only HEC can maintain that Otto holds a dispositional belief about the matter.³¹

Now a pressing problem arises for anyone who denies HEC, regarding how Otto can be said to know that MoMA is on 53rd street. Since friends and foes of HEC agree that Otto can form a belief with that content, it ought to be common ground that he can also come to know what he truly believes. After all, we can safely assume the reliability of his belief-producing process (relying on perception, for instance) and the absence of any undercutting or rebutting

³¹ For criticism of Chalmers & Clark's appeal to explanatory simplicity at this juncture see Rupert (2004: 420-424) and Sprevak (2009: 523-527).

defeaters. The key point is that Otto's knowledge involves dispositional belief rather than occurrent belief, but only HEC underwrites the claim that he dispositionally believes the target proposition.³² Otto (and Inga for that matter) knows a vast number of propositions, including that MoMA is on 53rd street, even when they are not consciously endorsed. What matters vis-à-vis knowledge is that these propositions be available to Otto (and Inga) for endorsement, which is exactly what a dispositional belief amounts to. In contrast, the Otto 2-step yields at best an occurrent belief in the target proposition. If Otto could only be ascribed knowledge on that basis, he would know next to nothing at any given time. For instance, his knowledge that MoMA is on 53rd street would begin and end with Otto consciously endorsing that proposition while consulting the notebook. When asleep, or when his mind goes blank, or when preoccupied with other mental activities he would be deprived of all notebook knowledge. Surely, Otto's knowledge could not be so discontinuous as to come and go with his explicit judgments on the matter.

What is more, if only HEC can account for Otto's knowledge of the notebook entries, because only HEC maintains that Otto holds the required dispositional beliefs, only HEC has the resources to causally explain behavioral differences by invoking either dispositional beliefs or knowledge. Thus, following Williamson (2000: 62-64), there may well be cases where knowledge can causally explain behavior no state short of knowledge can explain. Suppose Otto has been *Getterized*: someone with the same handwriting as Otto has accidentally written in the notebook that MoMA is located on 52nd street. On the basis of reading the notebook, Otto comes to believe by inference that the Museum is located in Midtown Manhattan. In this case, Otto's mere true belief can't explain why he stops looking for MoMA in Midtown Manhattan upon realizing the falsity of the lemma. In contrast, if Otto knows that MoMA is located in Midtown Manhattan, he will carry on looking even after his search on 52nd street proved unsuccessful. Of course, such cases are not entirely uncontroversial, but at least critics have assumed that dispositional beliefs were part of the explanans.³³ It's hard to see how the Otto 2-step can satisfactorily explain Otto's behavior in such cases when confined to occurrent belief in the target proposition. By contrast, since HEC allows for Otto to have memory knowledge, the broad or generic kind of memory this view upholds plays a distinctive causal-explanatory role with respect to Otto's behavior.

Here is the upshot. Clark has defended common-sense functionalism underlying the (METAPHYSICAL) PARITY PRINCIPLE by arguing that HEC offers distinctive explanatory

³² Rose and Schaffer (2013) provide arguments that knowledge entails dispositional belief.

³³ Magnus & Cohen (2003) is a case in point.

value to cognitive science.³⁴ We have defended such coarse-grained functionalism underlying our EPISTEMIC PARITY PRINCIPLE by arguing that HEC offers explanatory utility in epistemology: if the functional role of memory is set coarse enough to allow Otto to have dispositional beliefs, then not only can Otto be attributed memory knowledge, such knowledge can also be argued to play an irreducible role in causal explanations of his behavior³⁵.

§3.2 Otto and the Knowledge Condition

Even if we accept the foregoing—and it is granted that Otto exhibits a memorial process no less than Inga does—there obviously remains the final matter of explaining how Otto counts as possessing *knowledge* supported by memory no less than Inga. The worry, recalling here the Freya case, was that there is plausibly an ability condition on knowledge, and specifically a memorial ability condition on memory knowledge, which Inga, but not Otto, seems to satisfy.

Now, the turn-the-tables reply, in light of this worry, will advert to something like the following:

Credit Parity: To the extent that the correctness of a belief is creditable to a biological ability, it would also be creditable to what is (by reference to the parity principle) the extended functional analogue of that ability.

Credit Parity, in conjunction with Clark and Chalmers' parity principle, jointly yield the result that Otto's satisfies a memorial ability condition no less than Inga does. That is to say, if Inga's biological memory were included as a cognitive ability through an exercise of which her beliefs are true, but Otto's extended functional analogue were excluded as a cognitive ability in that sense, then the parity principle would be violated. In so far as the cases are functionally isomorphic, any such discrimination in favor of Inga in the attribution of cognitive abilities would give rise to bioprejudice against the case of Otto.

Interestingly though, at this point, friends of so-called *epistemic individualism* about cognitive abilities are going to reject Credit Parity. According to epistemic individualism, any such ability that plays a role in accounting for the truth of a belief in cases where an individual has knowledge has a biological or physical basis that is *wholly* resident in that individual. Thus,

³⁴ Sprevak (2009) argues that although the common-sense functionalism that Clark appeals to can save the Martian intuition, it does not sustain the fairly modest version of HEC that he puts forward. In fact, Sprevak claims any functionalism that preserves the Martian intuition supports a much stronger and overly permissive form of extended cognition/mind.

³⁵ We thank an anonymous referee for helpful remarks that have resulted in an expansion of this section.

those who endorse credit parity must adopt an *epistemic anti-individualist* position according to which the basis of any such ability is at most partially resident in the individual in question.

Epistemic individualists will, along with rejecting *Credit Parity* out of hand, request a further argument for why Otto's correctness should be down to his memorial ability. But it's not at all clear what further argument could be given, on behalf of Otto, *over and above* the kind of argument that could be given to demonstrate that the correctness of Inga's belief is appropriately down to her memorial cognitive ability, and then point out that his extended process (by the parity principle) shouldn't be treated any differently. After all, Otto's memorial process is instrumental in his representing the target belief in a way that doesn't vary from how Inga's process of consulting her biological memory is instrumental to her representing that belief. The turn-the-tables reply thus answers the challenge of explaining how Otto counts as possessing memory *knowledge* no less than Inga does by encouraging the epistemic individualist to look more closely at the cases of Otto and Inga, insisting that their relevant cognitive abilities be circumscribed on the basis of functional role as opposed to artificial bodily boundaries.

It's tempting to frame discussions of the ability condition on knowledge within an epistemic individualist framework, especially if the relevant cognitive abilities are conceived of along traditional reliabilist lines. For example, according to Goldman's early process reliabilism (1979; 1986), the cognitive belief-producing processes, whose reliability generates justification, are confined to the organic boundaries of the individual. Here's Goldman:

"A justified belief is, roughly speaking, one that results from cognitive operations that are, generally speaking, good or successful. But "cognitive" operations are most plausibly construed as operations of the cognitive faculties, i.e., "information-processing" equipment internal to the organism." (1979: 13)

"One thing we do not want to do is invoke factors external to the cognizer's psychology. The sorts of processes we're discussing are purely internal processes." (1986: 51)³⁶

So, although being reliable is what matters for the justificatory status of a belief, the cognitive processes that instantiate that property are restricted to events that are purely internal to the individual's psychology. Events outside the biological boundaries of the individual that are part of the causal ancestry of a belief count as mere inputs to the process that eventuated that belief,

³⁶ See also Alston (1995). Goldberg (2010: 36, 43-46) attributes to Goldman a view he calls "process individualism" according to which none of the processes that are implicated in the formation and sustainment of beliefs extend beyond the boundaries of the individual's head.

but they do not constitute a part of that process itself. Only cognitive events within the individual's organism do that. But then if, as on Goldman's view, all the relevant cognitive processes take place squarely within the individual's own brain or body, the corresponding cognitive abilities should naturally be understood as having biological or physical bases wholly resident in that individual³⁷. Against this background, it's not surprising, then, that the intuitions about creditability that are inherited from epistemic individualist thinking about cognitive ability are intuitions that lead us to require further argument for why the memorial ability condition is satisfied in the Otto case—after all, it won't be taken at face value in the way it is in the Inga case.

There is already a substantial amount of pressure on epistemic individualism in both social and mainstream epistemology, where it is increasingly argued that whether an individual has knowledge depends on epistemic factors beyond her cognitive agency that are incompatible with the epistemic individualist's thesis.³⁸ The most promising way to resolve the tension, we submit, is to abandon an entrenched conception of cognitive abilities as dispositional properties

³⁷ Almost categorically, epistemic internalists about justification have tacitly embraced epistemic individualism. Accessibilist internalists (Chisholm (1977); Ginet (1975) and Bonjour (1985, ch. 2)), for instance, simply take for granted that the process of reflection on the facts relevant to justifying a belief is a process grounded intracranially. (See here, especially, Ginet 1975: 38). Likewise, mentalists (e.g. Conee & Feldman (2004)) tacitly embrace epistemic individualism in virtue of supposing *ex ante* that the mental states on which epistemic justification supervenes are not (*a la* Clark & Chalmers 1998) extended mental states. See here Carter & Palermos (2014) and Carter, et al (2014) for further discussion.

³⁸ Interestingly, Sosa (2007: 93-98; 2011; 86-90) develops ideas that clearly move beyond epistemic individualism with respect to cognitive abilities. For instance, Sosa argues that testimonial knowledge involves complex competences that are socially seated, or seated in a group collectively. Also, we have elsewhere (Kallestrup & Pritchard 2012; 2013) provided a number of independent arguments for epistemic anti-individualism all of which assume the compatibility of this view with an ability condition on knowledge. More precisely, what these arguments show is that believing truly because of cognitive ability, where such ability is construed along epistemic individualistic lines, is not sufficient for knowledge. That is perfectly consistent with the epistemic anti-individualistic claim that believing truly because of cognitive ability is necessary for knowledge. And indeed when we argued in §2.3 that Freya, Memo and TrueTemp fall short of knowledge the claim was precisely that they failed to satisfy such necessary condition on knowledge, e.g. Memo lacks memory knowledge because his cognitive success is not through exercising a memorial ability on his part. Here is not the place to repeat these arguments in detail but a quick rundown is in order. Assume that individual S forms a true belief that p about her local environment through the exercise of her cognitive ability. As S's wider environment is conducive to knowledge, S comes to know that p. According to epistemic individualism, an internal, physical or biological, duplicate (S*) of S possesses the same cognitive ability as S. Moreover, assuming S* occupies an identical local environment, S* also forms a true belief that p through the exercise of her cognitive ability. However, unbeknownst to S*, S*'s wider environment contains knowledge-undermining features such that S* fails to know that p. It follows that what turns true belief into knowledge fails to supervene on internal properties, even on the assumption that the belief is true because of cognitive ability. Thanks to an anonymous referee for requesting further clarification on this point.

of individuals having biological or physical bases that are located exclusively inside their bodily boundaries. Once we conceive of Otto and his notebook as constituting an integrated cognitive system which gives rise to memory knowledge, the pertinent cognitive abilities can be viewed as dispositional properties of that system, which have bases that include features of both entities. Or so we claim, knowing full well that much hangs on how the details are cashed out in future work.

§4 Conclusion

We have raised here a novel tension for any attempt to reconcile HEC as it applies to knowledge-conducive processes with a plausible epistemology of memory. Our starting point is an individual Otto who purports to have knowledge supported by a memorial process that is both implemented non-biologically and extends beyond his bodily boundaries. If we suppose that Otto lacks knowledge due to merely parroting a prompt as in the case of Memo, then no reconciliation is forthcoming, and so this approach is a non-starter. Instead we explored the intriguing possibility that Otto possesses memory knowledge to the same extent that Inga does, except her memorial process is implemented biologically and located within her bodily boundaries. After all, parity of reasoning suggests those two cases receive identical epistemic treatments in light of their functional isomorphism. In that case, an explanation is needed of why the case of Otto displays neither relearning, nor triggering cues as in the case of Joggo. But even if that explanatory challenge can be met, a residual worry remains about the memorial cognitive ability an exercise of which accounts for Otto's putative knowledge. We recommend that the tables be turned. Once the phenomena of parroting prompts and triggering cues are carefully articulated as involving certain causal features, the case of Otto relevantly resembles that of Inga who indisputably displays neither. Moreover, when we abandon an entrenched conception of cognitive abilities along epistemic individualistic lines, the way is paved for a conception of the case of Otto as exercising an adequate cognitive ability necessary for memory knowledge.³⁹

References

- Alston, A. (1995). 'How to Think About Reliability', *Philosophical Topics* 23: 1-29.
- Atkinson, R.C., & Shiffrin, R.M. 'Human Memory: A Proposed System and its Control

³⁹ The authors would like to thank Sandy Goldberg, Emma C. Gordon, Rob Rupert, Orestis Palermos, Duncan Pritchard and an audience at *Eidyn* research centre for helpful comments on previous versions of this paper. Thanks also to the editor and an anonymous referee at *Philosophy and Phenomenological Research* who offered a number of very helpful suggestions. This article was written as part of the AHRC-funded 'Extended Knowledge' (#AH/J011908/1) research project that is hosted by the University of Edinburgh's *Eidyn* research centre.

- Processes. In: Spence, K.W., ed. *The Psychology of Learning and Motivation: Advances in Research and Theory*. New York: Academic Press, 1968. 89–195.
- Audi, R. (1994). ‘Dispositional Beliefs and Dispositions to Believe’, *Noûs* 28, 419-434.
- (1997). ‘The Place of Testimony in the Fabric of Knowledge and Justification’, *American Philosophical Quarterly* 34, 405-22.
- Ayer, A.J. (1956). *The Problem of Knowledge*. Harmondsworth.
- Bernecker, S. (2010). *Memory: A Philosophical Study*. Oxford: Oxford University Press.
- BonJour, L. (1985). *The structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Burge, T. (1986). ‘Individualism and psychology’, *Philosophical Review* 95: 3-45.
- (1993). ‘Content preservation’, *Philosophical Review* 102, 457–488.
- Chalmers, D. (2008). ‘Foreword to Andy Clark’s *Supersizing the Mind*’, In Andy Clark (ed.), *Supersizing the Mind*. Oxford: Oxford University Press.
- Carter, J.A. (2013). ‘Extended Cognition and Epistemic Luck’, *Synthese* 190 (18): 4201-4214.
- Carter, J.A. & Palermos, S.O. (2014). ‘Active Externalism and Epistemic Internalism’, *Erkenntnis*, in press.
- Carter, J.A. & Pritchard, D. (forthcoming). ‘Extended Entitlement’, forthcoming in *New Essays on Entitlement*, (eds.) P. Graham & N. Pedersen, (Oxford UP).
- Carter, A., J., Kallestrup, J., Palermos, S., O., Pritchard, D. (2014). ‘Varieties of Externalism’. *Philosophical Issues* 24: 63-109.
- Chisholm, R. (1977). *Theory of Knowledge*, 2nd ed., Englewood Cliffs: Prentice Hall.
- Clark, A. (2010). ‘Memento’s Revenge : The Extended Mind, Extended’. In R. Menary (ed.), *The Extended Mind*. MIT Press.
- (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.
- Clark, A. & David J. Chalmers (1998). ‘The Extended Mind’, *Analysis* 58: 7-19.
- Faulkner, P. (2011). *Knowledge on Trust*. Oxford: Oxford University Press.
- Ginet, C. (1975). *Knowledge, Perception, and Memory*. Dordrecht: D. Reidel Publishing Company.
- Goldberg, S. (2010). *Relying on Others*, Oxford: Oxford University Press.
- Goldman, A. (1979). ‘What is Justified Belief?’, In G. Pappas (ed.), *Justification and Knowledge*. Boston: D. Reidel, 1-25.
- (1986). *Epistemology and Cognition*, Harvard University Press.
- Greco, J. (2009). *Achieving Knowledge*. Cambridge: Cambridge University Press.
- Hall, N., 2000. ‘Causation and the Price of Transitivity’, *Journal of Philosophy* 97: 198-222.

- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Lewis, D. (1996). 'Elusive knowledge', *Australasian Journal of Philosophy* 74: 549-567.
- Kallestrup, J. & Pritchard, D. (2012). 'Robust Virtue Epistemology and Epistemic Anti-Individualism'. *Pacific Philosophical Quarterly* 93 (1):84-103.
- (2013). 'Robust Virtue Epistemology and Epistemic Dependence', in Henning, T. & Schweikard, D. Eds. *Knowledge, Virtue and Action: Putting Epistemic Virtues to Work*, Routledge.
- (2011). 'Robust Virtue Epistemology and Epistemic Twin Earth', *European Journal of Philosophy*, (Online First: doi: 10.1111/j.1468-0378.2011.00495.x).
- Kelp, C. (2013). 'Knowledge: The Safe-Apt View', *Australasian Journal of Philosophy* 93: 265-78.
- Lackey, J. (2008). *Learning From Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.
- (2009). 'Knowledge and Credit', *Philosophical Studies* 142: 27-42.
- (2007). 'Why We Don't Deserve Credit for Everything We Know' *Synthese* 158: 345-361.
- Lewis, D. (1973). 'Causation', *Journal of Philosophy* 70: 556-67.
- Magnus, P. D. & Cohen, J. (2003). 'Williamson on Knowledge and Psychological Explanation', *Philosophical Studies* 116: 37-52.
- Martin, C. B. & Deutscher, M. (1966). 'Remembering', *Philosophical Review* 75: 161-96.
- Matthen, M. (2010). Is Memory Preservation? *Philosophical Studies* 148: 3-14.
- McDowell, J. (1994). 'Knowledge By Hearsay', *Knowing from Words: Western and Indian Philosophical Analysis of Understanding and Testimony*, (eds.) B. K. Matilal & A. Chakrabarti, 195-224, Dordrecht, Holland: Kluwer.
- Menzies, P. 'Counterfactual Theories of Causation', *The Stanford Encyclopedia of Philosophy* (Spring 2014), Edward N. Zalta (ed.), forthcoming
URL=<<http://plato.stanford.edu/archives/spr2014/entries/causationconterfactua>>
- Moon, A. (2013). 'Remembering Entails Knowing', *Synthese* 190, 2717-2729.
- Miller, S. (2011). A Review of "Supersizing the Mind: Embodiment, Action, and Cognitive Extension", *World Futures* 66: 525-529.
- Pollock, J. (1970). *Knowledge and Justification*. Princeton University Press.
- Pritchard, D. (2012). 'Anti-Luck Virtue Epistemology'. *Journal of Philosophy* 109: 247-279.
- Putnam, Hilary (1975). The meaning of 'meaning'. *Minnesota Studies in the Philosophy of Science* 7: 131-193.
- Rose, D. & Schaffer, J. (2013). 'Knowledge entails Dispositional Belief', *Philosophical Studies* 166: 19-50.
- Rupert, R. (2004). 'Challenges to the Hypothesis of Extended Cognition', *Journal of Philosophy*

101: 389-428.

—— (2010). *Cognitive Systems and the Extended Mind*, Oxford: Oxford University Press.

Russell, B. (1921). *The Analysis of Mind*. London: George Allen and Unwin.

Ryle, G. (1949). *The Concept of Mind*. Hutchinson and Co.

Schaffer, J. (2000). 'Trumping Preemption', *Journal of Philosophy* 9: 165-181.

Senor, T. (2008). Epistemological problems of memory. *Stanford Encyclopedia of Philosophy*.

Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Vol. 1. Oxford: Oxford University Press.

—— (2011). *Knowing Full Well*. Princeton: Princeton University Press.

Sprevak, M. (2009). 'Extended Cognition and Functionalism', *Journal of Philosophy* 106, 503-527.

Theiner, G. (2011). *Res Cogitans Extensa: A Philosophical Defense of the Extended Mind Thesis*. Peter Lang.

Wheeler, M. (2010). 'In Defence of Extended Functionalism', In R. Menary (ed.), *The Extended Mind*. MIT Press.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.

Wilson, R.A. (2004). *Boundaries of the Mind: The Individual in the Fragile Sciences: Cognition*. Cambridge University Press.

—— (2008). 'The Mind Beyond Itself', In D. Sperber (ed.), *Metarepresentations: A Multidisciplinary Perspective*. Oxford: Oxford University Press.