Episteme, 12, 4 (2015) 459–478 © Cambridge University Press
doi:10.1017/epi.2015.30

EPISTEMOLOGY AND RADICALLY EXTENDED COGNITION

BENJAMIN JARVIS
benjamin_jarvis@alumni.brown.edu

ABSTRACT

This paper concerns the relationship between epistemology and radically extended cognition. Radically extended cognition (REC) – as advanced by Andy Clark and David Chalmers – is cognition that is partly located outside the biological boundaries of the cognizing subject. Epistemologists have begun to wonder whether REC has any consequences for theories of knowledge. For instance, while Duncan Pritchard suggests that REC might have implications for which virtue epistemology is acceptable, J. Adam Carter wonders whether REC threatens anti-luck epistemology. In this paper, I argue that the possibility of REC has no systematic consequences for theorizing in epistemology. I suggest an alternative relationship between the two: epistemology can play a role in diagnosing cases of REC. Thus, by establishing that entities partially located outside biological boundaries don't play certain epistemic roles, one can establish that they don't play the related cognitive roles either. I conclude the paper by illustrating this last point.

1. INTRODUCTION

The topic of this paper is the relationship between epistemology and radically extended cognition. *Radically extended cognition* (REC) – as advanced by Clark and Chalmers (1998) – is cognition that is partly located outside the biological boundaries of a cognizing subject.[1] The minimum requirement for REC is the partial extension beyond brain and body of cognitive processes; more extreme possibilities include the existence of radically extended cognitive states (e.g. beliefs) or radically extended cognitive character traits (e.g. standing belief-management mechanisms). Recently, philosophers have argued that REC is actual;[2] however, even critics allow that REC is at least possible.[3] Epistemologists have now begun to wonder whether REC has any consequences for theories of knowledge. For instance while Pritchard (2010) suggests that REC might have implications for which virtue epistemology is acceptable, Carter (2013) wonders whether REC threatens anti-luck epistemology. In this paper, I argue that the possibility of REC has no systematic consequences for theorizing in epistemology. I illustrate my point by considering the discussion of Pritchard (2010) and Carter (2013). I suggest an alternative relationship between epistemology and REC: epistemology can play a role in diagnosing cases of REC. By establishing

1 For some general overview, see Shapiro (2011).
2 Cf. Wilson (1994), Clark and Chalmers (1998), Rowlands (1999), Clark (2008), and Wheeler (2010).
3 Cf. Rupert (2009) and Adams and Aizawa (2010).

that entities partially located outside biological boundaries play certain epistemic roles, one can establish that they play cognitive roles as well. Similarly, by establishing that entities partially located outside biological boundaries don't play certain epistemic roles, one can establish that they don't play the related cognitive roles either.

Before I begin my arguments, some preliminary clarification may be helpful. To a first approximation, epistemology is the study of the correct principles of normative or value-laden assessment for a certain subspecies of the cognitive *qua* cognitive[4] – namely, doxastic processes, states, capacities, and mechanisms. Obviously, then, if some of these cognitive processes, states, capacities, and mechanisms are radically extended, they will be subject to epistemic assessment of the sort studied by epistemologists. Indeed, my view is not that there won't be possible cases of radically extended knowledge, rational graded belief, rational abilities, etc. There may be if radically extended (graded and non-graded) belief and radically extended belief-management mechanisms are possible. Moreover, it is uncontroversial that epistemic principles (broadly construed) will somehow cover possible cases of radically extended doxastic processes, states, capacities, and mechanisms. My contention is just that theories in epistemology that are formulated to be adequate for all possible cases of non-radically extended cognition will be adequate for all possible cases of radically extended cognition. In other words, theories in epistemology are, in a sense, immune to REC. The possibility of REC makes no difference to proper theorizing in epistemology, even if it has consequences for whether possible or actual units of epistemic assessment might be radically extended.

Because epistemological theorizing is immune to the possibility or actuality of REC, the failure of radically extended processes, states, and mechanisms to exhibit epistemic statuses could not be an independent threat to any particular *epistemology*. Instead, any such failure could only be a sign that these radically extended processes, states, and mechanisms are not *cognitive* units that would be subject to epistemic assessment. In other words, as already alluded, it is not as Pritchard (2010) and Carter (2013) have suggested – that REC has consequences for epistemology.[5] Rather, epistemic evaluation would appear to have consequences for the plausibility of REC. In effect, epistemic features are a "mark" of the doxastic family of the cognitive.[6] I illustrate this general point about the relationship by considering a case from Sutton *et al.* (2010) that purports to show that the memories of some people are radically extended into their partners. Towards the end of the paper, I contend that epistemic evaluation does not support this conclusion. In so doing, I demonstrate the merit of the methodology of using epistemology as a guide to diagnosing REC.

2. THE ROLE OF THE COGNITIVE IN EPISTEMIC THEORIZING

Unsurprisingly, theorizing in epistemology has been largely blind to REC (since prominent discussions of REC are recent). This is not to say that theorizing has been largely blind to

4 The "*qua* cognitive" qualification is necessary since doxastic phenomena can be subject to assessment that isn't epistemic. See, for example, Conee (1992) and Christensen (2004: §1.2).

5 Aizawa (2012) and Kelp (2013) make similar points restricted to virtue epistemology, but the point is fully general.

6 Adams and Aizawa (2012) emphasize the importance of establishing a "mark" of the cognitive for the purposes of diagnosing cases of REC.

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

cognitive kinds simpliciter, some of which might be instanced in radically extended ways. Indeed, the opposite is true. Epistemologists often take the cognitive to play a central role in their theories. The point is simply that epistemologists haven't formulated these theories in such a way that it would matter whether cognition is radically extended or not. In the past, theories of epistemic rationality and justification have been fleshed out in terms of cognitive kinds like belief, experience, perception, memory, inference, (cognitive) dispositions to believe, etc. Theories of knowledge have typically used the truth conditions of beliefs – a cognitive feature – to specify the factivity and anti-luck requirements for knowledge.[7] Also, it is commonly supposed that truth conditions of beliefs dictate something at least fairly central to their epistemology: a kind of objective success (truth) and a kind of objective failure (falsehood).[8] However, when epistemologists have invoked the cognitive in their theories, there's no evidence that it has been with any consideration for REC in particular. Certainly, invocations of the cognitive in epistemology haven't included any restrictions as to whether the cognitive in question is radically extended or not. For instance, epistemologists haven't tended to say things like "Beliefs can be properly inferred from other beliefs unless the latter beliefs are radically extended" or "Knowledge is factive and safe unless it is radically extended knowledge" or "If ducks quack, a belief that ducks quack is objectively successful even if not justified – unless, of course, the belief is radically extended."

It is natural to be suspicious of blanket generalizations about practice in epistemology, but it is fairly easy to check that, at least up until very recently, the idea of REC hasn't made an appearance in standard anthologies in ("analytic") epistemology.[9] Nevertheless, given the importance of the point for latter discussion, it may be worth surveying some examples to see that, in fact, prominent theories in epistemology are formulated so as to be indifferent to the possibility of REC. I offer five:

(1) Consider the mentalism of Feldman and Conee (2001): "The justificatory status of a person's doxastic attitudes strongly supervenes on the person's occurrent and dispositional mental states, events, and conditions." This theory of epistemic justification says precisely that justification depends only on a subject's total mental condition in the present. While this principle doesn't explain how justification depends on a subject's total mental condition, it doesn't suggest that radical extension should matter. In fact, the principle seems to entail that it shouldn't matter. According to the principle, the *only* relevant consideration is which mental states, events, and conditions are presently instanced by the person. Consequently, it would appear *not to be relevant* whether these mental states, events, and conditions are radically extended or not.

Other theories of justification in the literature are also indifferent to the possibility of REC:

(2) Feldman and Conee (2001) contrast their view with *accessibilism*.[10] To a first approximation, accessibilism says that a belief is justified to the extent that other (specially) accessible mental states lend it support. Note that (special) accessibility is a *cognitive* feature. Depending on how (special)

7 See, for example, Pritchard (2005).

8 On the importance of truth as objective success, see, for instance, Lynch (2004, 2009). On the importance of falsehood as objective failure, see Whiting (2012, 2013).

9 See, for example, Sosa et al. (2008).

10 They provide references although as a point of comparison, one might look at the works of Richard Foley, e.g. Foley (1993).

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

accessibility is understood, it may be that (specially) accessible mental states must be wholly located within biological organisms – so that a certain kind of radically extended cognition is impossible. However, accessibilism itself *does not differentiate* between radically extended accessibility (if any exists) and non-radically extended accessibility.

(3) A salient rival to mentalism is Goldman's (1979) *process reliabilism*. To a first approximation, reliabilism says that a belief token is epistemically justified if and only if it is the deliverance of a token process that is type-reliable. A central difficulty of this theory is accounting for what it is for a token process to be type-reliable, since any token process will fall under any number of types, some of which will be reliable, some of which won't.[11] Which of these types, we might wonder, is the type that matters for determinations of type-reliability? If reliabilism is to be completed in such a way as to render verdicts about particular cases, a principled answer to this question is required. It seems fairly plausible that the type that matters for determinations of reliability will be partly cognitive in nature. For instance, reliabilists seem to think that beliefs formed by visual perception will turn out to be epistemically justified on this theory because cognitive processes supporting vision are reliable. Goldman (1979) suggests the relevant types are individuated *only* by "cognitive" events. He also suggests that these will be "nervous system" events. However, this further suggestion seems to arise from an unwitting conflation of the boundaries of the cognitive system with boundaries of the nervous system rather than a deliberate commitment to importance of the latter for the theory of justification. Once the distinction between cognitive and neurological boundaries becomes clear, it also becomes fairly clear that the relevant boundaries for typing processes are cognitive. If so, epistemic justification of beliefs will not ultimately turn on whether the beliefs were the deliverance of token processes that fall under the type qualifying as radically extended cognition. Types of processes relevant for epistemic justification will be determined by what (*qua* cognitive) happens within the cognitive system whether or not it happens in a radically extended way.

The neutrality of epistemology with respect to REC does not end with theories of justification:

(4) Consider Williamson's (2002) theory of *propositional evidence*. According to this theory, a proposition is among one's evidence (and has epistemic probability equal to one) if and only if one knows the proposition. Since knowledge entails belief, it would appear that this theory requires that a subject believe any proposition that is evidence for her. In other words, there are cognitive requirements for evidence. However, there are no radically extended cognitive requirements for evidence. Nor does being believed in a radically extended way disqualify a proposition as evidence. REC simply makes no difference when it comes to evidence. (This remains true even if we add Williamson's claim that knowledge is a distinctive cognitive kind. Knowledge of external facts isn't radically extended per se. Instances of knowing are token identical to instances of believing, even though what qualifies a belief as knowledge is, inter alia, a fact that very typically is external to the subject. Thus, knowledge might well be located in the head even if what makes it knowledge isn't.)

(5) Finally (and perhaps most saliently for our purposes), consider *robust virtue theories of knowledge*.[12] According to a robust virtue theory of knowledge, an agent's belief qualifies as knowledge if and only if it is an objective success (i.e. a true belief) and this ongoing success is attributable (as such) to the agent's exercising of her rational abilities. This biconditional allows that knowledge might be explained in terms of objective success and rational abilities (on the condition that the left half of the biconditional functions as the explanandum while the right half functions

11 *Cf.* Conee and Feldman (1998) and Comesaña (2006).

12 Greco (2003, 2010), Sosa (2003, 2007), and Zagzebski (1994).

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

as the explanans). Consequently, it might appear that REC matters for the theory of knowledge if it matters for the theory of rational abilities. However, the epistemology of rational abilities appears to be *neutral* with regards to the possibility of REC. To a first approximation, rational abilities are (stable) cognitive character traits that are intellectually beneficial (e.g. promote objective successes or diminish objective failures). This theory of rational abilities allows for radically extended rational abilities if and only if they are radically extended cognitive character traits. Yet, there is no indication that qualifying as rational abilities (rather than mere belief-management mechanisms) turns on whether or not cognitive character traits are radically extended. What matters is simply being a cognitive character trait and being intellectually beneficial (in the right way, whatever that turns out to be). Moreover, being a cognitive character trait is a cognitive feature to be theorized on not by epistemologists qua epistemologists, but by philosophers of cognition and cognitive scientists instead.

While an exhaustive survey is not possible, (1)–(5) provide enough of a sampling to support the claim that theorizing in (at least "analytic") contemporary epistemology hasn't generally allocated any significance to *radically extended* cognition. The radical extension of cognition hasn't been taken to explain the presence or absence of epistemic properties.

The question now arises, however, whether this blindness to REC is problematic. Epistemic statuses are, like all normative or value-laden statuses, subject to supervenience constraints.[13] There cannot be differences in epistemic status without underlying (mental, functional, or material) differences. Epistemic assessments do not cohere if they attribute epistemic differences without compensatory differences of other kinds. One might propose that the difference between cognition that is and isn't radically extended just is the kind of difference that could ground epistemic differences (even though epistemologists haven't taken them to). The distinction between cognition that is and isn't radically extended might be theoretically important for epistemology in just this way. But, how plausible is this proposal?

Not very, I think. While I don't have any knockdown arguments, it strikes me that the burden of proof rests on any proponents who, in effect, favor special clauses in epistemological theories to cover cases of REC. The basic line among believers in REC is that the distinction between those processes, states, and character that are radically extended and those that aren't doesn't matter for the philosophy of cognition.[14] However, if this distinction doesn't matter for cognition, it's hard to see why it should matter for the assessment of cognition as such – and assessment of certain (doxastic) kinds of cognitive as such is, of course, precisely what epistemic assessment is.

Thomas Jefferson – or more specifically, how he was – is arguably relevant to the epistemic assessment of beliefs about Thomas Jefferson (e.g. whether they meet epistemic standards for success) at least partly as a consequence of the fact that beliefs about Thomas Jefferson only qualify as beliefs of this kind due to their relation with Thomas Jefferson.[15] This illustrates the general principle that if *Y* is relevant to whether *X* is cognitive in a particular way, then *Y* may well be relevant to whether *X* is, in its role as cognitive in that particular way, good, bad, etc. However, a qualified converse of this principle is also plausible. If *Y* isn't relevant to whether *X* is cognitive in a particular way, it is not at all clear why *Y* should be relevant to whether *X* is, in its role as cognitive in that particular way, good, bad, etc. By way of example, beliefs about Australia do not

13 Hare (1952).

14 Wheeler (2010) provides a very clear characterization of this idea.

15 The latter follows from Putnam–Burge semantic externalism (Putnam 1975; Burge 2007, 2010).

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

qualify as beliefs of this kind due to their relation with Thomas Jefferson. In specific circumstances, facts about Thomas Jefferson might be relevant as such to epistemic assessment of beliefs about Australia, but a special explanation is presumably required as to why – e.g. an explanation of appropriate relations between Thomas Jefferson and Australia. Certainly, it is not at all clear why facts about Thomas Jefferson would be *generally* relevant as such to epistemic assessment of beliefs about Australia. Consider a further analogy. It is not at all surprising that relations between Jim and Jane matter for the evaluation of Jim *qua* father if Jim is a father in virtue of his relation to Jane (because Jim is the father of Jane). But, if Jim is not a father in virtue of his relation to Jane, it is not at all clear why the relations between Jim and Jane should matter for assessing Jim as a father. Of course, they might matter, but it is reasonable for us to expect an explanation of the relevance. Similarly, in the case of REC, being radically extended isn't (so we are told) relevant to being cognitive per se, so it is not at all clear why radical extension should generally be relevant to epistemic assessment of the cognitive, i.e. assessment of the cognitive *as such*.

Let me stress further that I am not making a definitive claim about the relevance of REC to epistemology. I am merely indicating that it is reasonable to expect an explanation as to why radical extension should be relevant per se to the epistemic when it isn't to the cognitive. In the absence of this explanation, I contend it is appropriate to proceed on the assumption that the radically extended distinction – i.e. the distinction between those processes, states, and character that are radically extended and those that aren't – isn't one that should appear in theories in epistemology. If this assumption is true, then the blindness that epistemologists have shown to the possibility of REC isn't obviously problematic. Epistemologists haven't been saying things like "Knowledge is factive and safe unless it is radically extended knowledge," but they shouldn't have been. Rather, they should have been formulating their theories using cognitive kinds that are neutral with respect to radical extension – just as they have.

It's worth noting that any defense of epistemic parity principles[16] *ipso facto* goes to support the idea that (at least certain) epistemological theories should not include carve-out clauses for REC. For example, Carter and Pritchard (forthcoming) defend the following principle:

> *Epistemic Parity Principle**: For any agent *S* and true proposition *p*, if *S* comes to believe that *p* by a process which, were it to go on in the head, we would have no hesitation in ascribing defeasible warrant to *S*'s belief that *p*, then *S* has defeasible warrant for her belief that *p*.

Any theory of defeasible warrant that included a carve-out clause for radically extended cognitive processes would violate this principle because it would allow for a case where the presence of defeasible warrant turned on whether the REC carve-out clause applied. So, if this epistemic parity principle is true, theories of defeasible warrant should not include carve-out clauses for REC of this sort.

Even still, the failure of certain epistemic parity principles need not undermine the main claim of this section. For instance, Carter and Pritchard (forthcoming) argue against this principle:

16 Cf. Carter (2013) and Carter and Pritchard (forthcoming). Epistemic parity principles mimic the Parity Principle of Clark and Chalmers (1998).

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

Epistemic Parity Principle: For any agent *S* and true proposition *p*, if *S* comes to believe that *p* by a process which, were it to go on in the head, we would have no hesitation in ascribing knowledge that *p* to *S*, then *S* knows that *p*.

However, this principle fails to hold for general considerations that do not stem from the possibility of REC. *Any* difference in feature, *F*, between two beliefs, B and B*, can make for a difference in knowledge between B and B* if the agents are in a position to rationally suspect that beliefs with feature *F* are generally or locally unreliable. Thus, the fact that a belief was formed by a REC process could, in principle, go towards explaining why the belief isn't knowledge because an agent might be in a position to rationally suspect that beliefs formed by REC process are generally unreliable – or, at least, unreliable in the present circumstances. Nevertheless, no special REC carve-out clauses are required in the theory of knowledge since this last point can be fully generalized so as to abstract away from REC in particular. Consequently, the main claim of this section – that radical extension is not relevant *per se* to the epistemic – goes untouched.

3. RADICALLY EXTENDED COGNITION IN COUNTEREXAMPLES

If epistemologists shouldn't be carving out exceptions for the epistemic assessment of REC, it becomes hard to see how REC could be an important new source of counterexamples to epistemological theories. In this section, I argue this is exactly right. I illustrate my argument by considering two potential counterexamples from the philosophical literature.

First, though, let me offer a principled explanation as to why REC couldn't be a new source of counterexamples in epistemology. As I pointed out in the immediately preceding section, one of the guiding ideas of the REC literature is that cognitive roles occupied by internal biological entities could, in principle, be occupied by radically extended entities instead.[17] However, if this guiding idea is correct, one might also expect that cognitive roles occupied by radically extended entities could, in principle, be occupied by internal biological entities. Consider, then, a hypothetical counterexample to an epistemological theory that invokes REC. It would appear that there should be a structurally analogous case in which any radically extended entities playing cognitive roles in the original counterexample are replaced by internal biological entities instead. Assuming the epistemological theory, while invoking various cognitive distinctions, is nonetheless insensitive to REC per se, the structurally analogous case should be just as effective as a counterexample. In other words, epistemological theories are only prone to counterexamples with REC, if they are prone to counterexamples without it.

It will help to illustrate how one might construct a counterexample without REC from an example with. Consider first a potential counterexample from Pritchard (2010). He considers a variation on the famous Otto example of Clark and Chalmers (1998). In this variation, Otto's (allegedly) radically extended memory system – in the form of a notebook – is set up for him by his wife. This (allegedly) radically extended memory system delivers true beliefs, but Pritchard suggests these true beliefs are primarily creditable to Otto's wife. This, he judges, may create a problem for the robust virtue theories of knowledge described earlier. Otto knows, but these true beliefs aren't attributable as such to the

17 Again, for a clear statement of this idea, see Wheeler (2010).

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

exercise of his (radically extended) rational abilities (or so Pritchard alleges) because they are attributable to the wife's designs and execution instead.

For our purposes, it doesn't matter whether Pritchard's case is in fact a (successful) counterexample (even though I think it isn't).[18] Rather, the issue is whether there is a possible case without REC that would be equally effective as a counterexample.[19] I would argue any problem for robust virtue theories of knowledge is equally well created by a case in which the memory system clearly isn't radically extended. For instance, imagine Otto's wife designs a miraculous pharmaceutical drug. In combination with a very strict diet and exercise regimen, which the wife dutifully enforces, this drug bio-enhances (in a distinctively non-natural way) Otto's memory so that it functions as well as the radically extended memory system in the other case. If true beliefs delivered by the radically extended memory system are primarily creditable to Otto's wife, then true beliefs delivered by this bio-enhanced memory system are primarily creditable to Otto's wife too. Notice that in generating a new example, I have simply replaced radically extended entities playing cognitive roles with internal biological entities playing the same roles. As a result, the two cases are structurally analogous in the relevant respects so that they should function equally well as counterexamples. Because in the latter case the memory character traits are not radically extended, it becomes clear that REC is not an essential ingredient for constructing a counterexample.

The same kind of procedure also works for a second potential counterexample from Carter (2013). Carter (2013) is interested in the thesis that knowledge is incompatible with environmental luck.[20] By way of background, environmental luck might be illustrated by a case in which you form (as it happens) a true belief that you have a five dollar bill in your wallet from having seen it.[21] Ordinarily, the visual experience of the bill might well suffice for knowledge. However, in this particular case, there are a large number of counterfeit five dollar bills in the local area due to a recent visit by North Korean spies. These counterfeit five dollar bills could only be detected by special machines at the Federal Reserve. In this case, it is lucky that you have a true belief that you have a five dollar bill in your wallet even though there is nothing defective about the process by which you form this belief. Lack of defect in the belief-forming process is, in fact, a distinguishing feature of environmental luck. A true belief is environmentally lucky if and only if it is formed in a way that (1) is generally effective at yielding true beliefs, but (2) given local environmental conditions, is not locally effective at yielding true beliefs, but (3) it is rational to believe that the local environmental conditions are not special in this respect (even though, in fact, they are).

It is commonly held that knowledge is incompatible with environmental luck.[22] Carter (2013) worries that this case might be problematic for the commonly held thesis:

> JOKESTER: 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

18 Even if the success is attributable to the wife, this doesn't preclude the success also being attributable to Otto. See Jarvis (2013) for some discussion. Cf. Nozick (1974: 129–30) on responsibility.

19 Cf. Aizawa (2012) and Kelp (2013).

20 On the contrast between environmental and intervening veritic luck, see Pritchard (2005).

21 The classic environmental luck case is the barn façade case from Goldman (1976).

22 Carter (2013: fn. 7) provides a number of references.

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

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.

Ignoring the possibility of REC, the JOKESTER case looks like a relatively clear case of environmental luck. There is nothing intrinsically defective with the process by which Otto comes to have an occurrent belief that his appointment with his doctor is at noon. Ordinarily, relying on his notebook – presumably, a kind of self-testimony – does not leave him with false occurrent beliefs; generally, this procedure is effective. However, given the unusual presence of the jokester, this procedure is locally ineffective. Consequently, it looks as if the conditions for environmental luck are met. If knowledge is incompatible with environmental luck, then it seems that Otto does not know in the JOKESTER case.

This last conclusion isn't so clearly right, however, if we suppose that Otto's memory system is radically extended into the notebook in the way that Clark and Chalmers (1998) suggest. Relying on his notebook is not merely a kind of self-testimony for Otto (if it is self-testimony at all); it is ordinarily a kind of *remembering*. Moreover, the jokester is literally tampering with Otto's memory by fudging the appointment times listed in the notebook. There continues to be a sense in which the belief about the doctor's appointment is lucky; it is lucky that Otto's memory vis-à-vis his doctor's appointment hasn't suffered interference from the jokester. Very easily the jokester might not have overlooked the doctor's appointment listed in the notebook. However, it is not obvious that *this* kind of luck is incompatible with knowledge. Perhaps, Otto knows in the JOKESTER case after all.

Carter (2013) worries that, if JOKESTER is understood to be a case of REC, it threatens to undermine the thesis that knowledge is incompatible with environmental luck. JOKESTER, he suggests, continues to be similar to paradigmatic cases of environmental luck – for example, the counterfeit money case that I introduced earlier. Thus, Otto should not know in JOKESTER, if subjects don't know in paradigmatic cases of environmental luck. As it happens, I find this reasoning problematic: however similar it might be to paradigmatic cases of environmental luck, JOKESTER only has to be different in one relevant respect to merit a different kind of verdict.[23] Moreover, as I highlighted earlier, cognitive distinctions are generally relevant to proper epistemic assessment. If Otto's memory system is radically extended into his notebook, then the problem (if any) in JOKESTER isn't that local environmental conditions aren't favorable for Otto's acquiring a true belief about the doctor's appointment by ordinary means. Rather, the problem appears to be located *inside Otto's mind* with his memory storage. In other words, JOKESTER is not a case of environmental luck; it is (at least superficially) similar to cases of environmental luck, but cognitively different. I don't see why we should expect that this cognitive difference won't matter to determinations of knowledge.

Nevertheless, we might suppose Carter (2013) is right and JOKESTER does threaten the thesis that knowledge is incompatible with environmental luck. If so, then my view is that there should be some case that doesn't invoke REC, but threatens the thesis equally well. In fact, Carter (2013) suggests a possibility,[24] but a more closely analogous case might be this one:

23 Carter is sympathetic to this complaint in his current (unpublished) work.

24 Carter's (2013) case:

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

BIO-JOKESTER: Otto* (without Alzheimer's) has a normally functioning biological memory, which he consults to determine when his doctor's appointment was today, and finds the correct time, noon, is what he (at least seems) to remember. Unbeknownst to Otto*, his memory has been systematically altered by a jokester, who used pharmaceuticals and subliminal suggestion to plant false memories about his other appointments that day, making Otto* inclined to believe that they are an hour earlier than he used to believe. The jokester, however, overlooked the doctor's appointment, leaving the original and correct time registered in Otto*'s memory.

Short of skepticism about whether there really is REC in cases like JOKESTER, I don't see any reason to think that there should be a difference in verdict between JOKESTER and BIO-JOKESTER. If one grants that Otto's memory is radically extended in JOKESTER, then I don't see any basis for thinking that Otto fails to know the time of the doctor's appointment while Otto* knows or vice versa. Moreover, if the similarity between the JOKESTER case and cases of environmental luck does, in fact, indicate that verdicts about knowledge in those cases should match, then the same should be said of the BIO-JOKESTER case and cases of environmental luck. The BIO-JOKESTER case is similar to cases of environmental luck in all relevant respects – or, at least, it is if we assume that Otto and Otto* are cognitively alike. (Note that the conclusion from the previous section implies that the obvious difference between Otto and Otto* on this assumption – i.e. how far cognition extends – is *not* relevant to epistemic determinations.) Thus, if JOKESTER provides reason to worry that knowledge is compatible with environmental luck, so does BIO-JOKESTER. The upshot is that Carter (2013) is wrong to suggest that Otto-style REC presents a special problem for the thesis that knowledge is incompatible with environmental luck. Any such problem for this thesis exists *independently*.

I began this section with a *principled* explanation of why we ought to expect that theories prone to counterexamples involving REC should also be prone to counterexamples without REC. The cases of Pritchard (2010) and Carter (2013) have provided me the opportunity to illustrate how to manufacture potential counterexamples without REC from potential counterexamples with it. To do so, one simply replaces instances of REC with equivalent instances of non-radically extended cognition, *mutatis mutandis*. I suggested in the previous section that theories in epistemology should not have special carve-out clauses for REC. If they do not, potential counterexamples manufactured by the illustrated method will serve equally well as the original counterexamples with REC.

It follows that any epistemological theory that is materially adequate when restricted to cases without REC will be materially adequate tout court. (If it weren't materially adequate tout court, then there would be a counterexample. There are two possibilities – (a) that the counterexample invokes REC and (b) that it doesn't. If it does, then another counterexample can be manufactured that doesn't invoke REC. Either way, there is a counterexample that doesn't invoke REC. This proves the contrapositive.)

Case 2: Otto* (without Alzheimer's) has a normally functioning biological memory, which he relies on to organize 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.

There are at least two problems with the suggestion that this case is an analogue of JOKESTER. First, forgetting and misremembering are normal happenstances (unlike jokesters messing with notebooks). Second, the case doesn't imply that the integrity of the doctor's appointment memory was in peril.

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

Thus, for the purposes of theorizing in epistemology, one can simply ignore the possibility of REC and focus on clear cases of non-radically extended cognition. Or, put another way, the relationship between REC and epistemology is *not* that the former has consequences for the latter. Unlike, for example, the possibility of epistemic luck, the possibility of REC makes no difference to epistemology.

In closing this section, it's worth briefly addressing a particularly dramatic way in which cases of REC might alter theorizing in epistemology – that of inspiring a paradigm shift. For instance, Vaesen (2014) suggests that reflecting on REC might result in a paradigm shift from "analytic" epistemology to epistemology of a more pragmatist bend. In particular, by thinking about how Otto's notebook is interacting with biological-situated states and mechanisms to get him where he's going, we might be better positioned to appreciate the pragmatic "steering" role of doxastic elements (including doxastic activities) as well as the importance of considering the doxastic system and its dynamic activities as a whole rather than individual beliefs in isolation.[25]

For the record, I don't think that "analytic" epistemology has any particular difficulty in adequately addressing the "steering" role of the doxastic,[26] nor does it incorporate an unduly narrow focus on the evaluating individual beliefs,[27] but, for present purposes, we can put those issues aside. Nothing put forward in this paper commits me to thinking that REC cannot *inspire* better epistemological theorizing. Rather, the relevant point is that cases of REC would *not be needed to justify* theoretical revisions in epistemological – including revisions as dramatic as a paradigm shift. If Otto's notebook is especially different from biological memory in its steering role, connectivity, or general dynamic activity, then motivation for judging it to be a case of REC akin to biological memory is lost. I assume, then, for the sake of argument, that biological memory exhibits similar properties. It follows that if Otto's notebook is indicative of the pragmatic "steering" role of doxastic elements or the importance of considering the doxastic system and its dynamic activities as a whole, then the correlative biological memory is as well. So, even if the target is as large as 'analytic' epistemology itself, using a case of REC seems inessential.[28]

25 Vaesen (2014: 435).

26 In fact, the demise of behaviorism taught us that individual beliefs don't have a specific role in steering action; the way that a belief steers action depends on other beliefs and desires. Truth-centric epistemic evaluation is pertinent to understanding the pragmatic credentials of a belief since (1) truth propagates through a system of appropriately linked beliefs and (2) bringing it about that p will (likely) be a successful action if based on a desire – with content q – that is good and an instrumental belief – with content *(likely)* q *if* p – that is true. Cf. Horwich (2006). (Any role of truth in ensuring successful action is, of course, compatible with thinking that truth and other related properties do not exhaust or wholly determine the pragmatic credentials of a belief.)

27 The orthodoxy since Quine (1953) is confirmation holism, which straightforwardly entails that beliefs don't obtain the status of being justified, rational, or even knowledge in isolation from other elements in the doxastic system.

28 It may be worth pointing out that if the aim is to emphasize the importance and "practical" aspects of dynamic mind-environment activity for epistemology, then REC seems particular unsuitable for making the point. When REC is actual, the boundaries of the mind are further extended – meaning there's less dynamic cross-border activity, not more.

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

4. RADICALLY EXTENDED COGNITION AND THE PRACTICE OF EPISTEMIC EVALUATION

Nevertheless, as indicated in the introduction, REC does have some consequences for the distribution of epistemic properties in the world. In other words, although epistemological *theorizing* should be unaffected by REC, *the practice of epistemic evaluation* should be affected. The reason is that proper epistemic evaluation requires sensitivity to whether something is a proper unit of epistemic evaluation. One should not evaluate an accidental, unintended tripping as irrational merely because it fails to satisfy the preferences of the person who tripped; tripping is not usually an action and so is not properly evaluated as such. Unintended tripping isn't a proper unit of action assessment. Similarly, one should not evaluate some scribbling in a notebook as an irrational state of mind, if it isn't a state of mind at all. If the scribbling in the notebook isn't a belief, then it shouldn't be epistemically evaluated as a belief would be. On the other hand, if the scribbling in the notebook is a belief, then it should be epistemically evaluated as a belief would be. Judgments that token processes, states, capacities, and mechanisms are instances of some doxastic-related variety of the cognitive commit one to epistemically evaluating them as such. Moreover, epistemic evaluation of token processes, states, capacities, and mechanisms commits one to judgments that these tokens are cognitive. In addition, *the way* in which one epistemically evaluates something typically commits one to judging it to be a particular kind of cognitive. For purposes of illustration, consider the contrast between intuition and belief. An intuition and a belief that bodiless pain is possible each clash in different ways with a belief that bodiless pain isn't possible. Having beliefs with contradictory contents is markedly worse, from an epistemic standpoint, than having an intuition and a belief with contradictory contents.[29] Consequently, epistemically evaluating a mental state token as very strongly clashing with another mental state token with contradictory content gives rise to certain commitment. One is committed to judging that both tokens are beliefs rather than intuitions.

The fact that REC has some consequences for the distribution of epistemic properties in the world is obviously related to the very tight relationship between judging something to be cognitive in a particular way and epistemically evaluating it. However, this same very tight relationship is just as related to the fact that the distribution of epistemic properties in the world has consequences for REC. (Consequence, in the relevant sense, is a metaphysical relation, but not an explanatory one. Consequence is not asymmetrical; the distribution of A-properties may have consequences for the distribution of B-properties even if the reverse also holds true.) Moreover, as I pointed out, there are, in fact, *two way* rational commitments between judging-as-cognitive and evaluating epistemically, so one might expect two-way consequences between the cognitive and epistemic domains. Notably, it is not clear that our grasp on what is cognitive should be in any way prior to our grasp on how epistemic properties in the world will be distributed.[30] Indeed, even if the cognitive is *metaphysically* prior to the epistemic,[31] *our sensitivity* to the

29 Cf. Koksvik (unpublished).

30 Consider that we seem to have a much better grasp on epistemic roles than on cognitive roles. From the armchair, it may be much easier to explain, e.g., what would count as evidence for the belief that the Federal Reserve's loose monetary policy is causing economic expansion than it is to offer either empirical generalizations involving this type of belief or various causal conditions that are necessary/sufficient for a mental state to qualify as being a belief of this type.

31 This order of metaphysical priority is denied by normativists, e.g. Zangwill (2005, 2010).

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

cognitive may be mediated in a variety of ways by our sensitivity to the epistemic. As a result, it may be that the best way to weigh the merits of any particular claim of REC is by turning to the epistemic domain. In effect, one should be able to defend a judgment that some radically extended token is of a particular kind of cognitive by defending the epistemic evaluation of it *as such*.

For example, in other work, I have defended the view that certain radically extended belief-forming processes – pen and paper computations of large arithmetic sums – are cognitive processes by defending the view that these processes are subject to the corresponding kind of epistemic evaluation.[32] In effect, aspects of these radically extended processes determine whether resulting belief tokens are held because they are evidentially supported, a property these processes would have only if they were part of a distinctly cognitive performance. In fact, one should also be able to defend a judgment that some radically extended token *isn't* of a particular kind of cognitive by defending the view that it *shouldn't* be epistemically evaluated in the requisite way. I will turn to this strategy shortly when I argue that certain radically extended processes are not instances of remembering because they are subject to the same kind of epistemic evaluation as genuine remembering is.

Before I do so, however, it is important to emphasize that the effectiveness of these strategies for diagnosing REC turns on the results from §§2–3. Suppose that a radically extended token is alleged to be of a certain cognitive kind, yet this token doesn't exhibit the same epistemic features that an otherwise analogous non-radically extended token of this cognitive kind would. Because REC could not be a genuinely new source of counterexamples in epistemology, this could only mean that the allegation is false. However, if, *pace* my earlier discussion, REC is a source of counterexamples, then the allegation might be true anyway. The epistemic difference might be explained by the difference between radically extended and non-radically extended cognitive kinds. In other words, noting this epistemic difference might provide, for example, the rational impetus to revise an epistemological theory so as to include a carve-out clause for REC. If epistemic difference might be attributable to radical versus non-radical extension, epistemic difference need not undermine cognitive sameness. Thus, it is precisely because epistemological theorizing is immune to the possibility or actuality of REC that the presence of epistemic features can be used, in principle, to straightforwardly diagnose the presence of the cognitive. In other words, the methodology of using the epistemic as a "mark of the cognitive" crucially depends on the neutrality – claimed in §§2–3 – of epistemology with respect to radically versus non-radically extended cognition.

5. AN ALLEGED CASE OF RADICALLY EXTENDED COGNITION

Whether, *in practice*, we can use the presence of epistemic features to diagnose the cognitive depends on whether, in fact, we might have a prior awareness of the epistemic features. It is a theoretical possibility that any awareness of epistemic features comes via our overt knowledge of the cognitive. If so, we're not going to learn about the distribution of cognitive features in the world by learning about the distribution of epistemic features since we learn about the latter by already having learned about the former. I've proposed

32 Jarvis (2014).

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

the methodology of looking for the cognitive in the epistemic. But, for this methodology to pan out, we need to be more competent at making explicit epistemic evaluations than we are at judging whether something falls under a particular cognitive kind.

As it happens, we are, at least sometimes, better attuned to normative properties than non-normative one. Famously, subjects do better at identifying potential violations of "If a person is drinking, the person is over the age of twenty-five" than potential counterexamples to "If a card shows an even number on one face, then its opposite face is red."[33] A possible explanation is that detecting cheaters is of far more importance than uncovering counterexamples to abstract generalizations. For similar reasons, it seems plausible enough that we might be more competent at making epistemic evaluations than we are at judging whether something falls under a particular cognitive kind. The epistemic credentials of a person's beliefs or character traits may be, in many cases, more relevant to how to treat the person than the particulars about what exactly in the vicinity of the person is cognitive and in what way. Perhaps, then, due to its practical import, human beings have developed a more intuitive grasp of the epistemic than they have of the cognitive.

Indeed, it is enough for my purposes that it is theoretically possible that we are better at making epistemic evaluations than judgments about the cognitive, and, also, that there are possible explanations for why this might be so. When it comes to justifying the methodology of looking for the cognitive in the epistemic, it wouldn't be wise to rest much on a just-so story. The best way to support this proposed methodology is rather to show how it works in a particular case. I propose to consider a case of remembering – recorded by Sutton *et al.* (2010) – that is advanced as a potential instance of REC.[34] The case involves a husband and wife remembering in tandem the beginning of their relationship:

Husband: No, I asked her out that night, but she said she couldn't go.
Wife: No, that's right.
H: So then I started to pester her the next week.
W: You did, you turned up after my [classes].
H: [Cooking classes].
W: On Monday night.
H: That'd be it.
W: And took me for coffee.
H: Yes, the next Monday night.
W: And impressed me.
H: Yes.

The question I will consider is whether the husband's process of remembering, say, that it was Monday night when the pestering began is radically extended into the wife (and vice versa). (There is at least one other theoretical possibility for REC here: the husband and wife could collectively constitute another distinct epistemic agent who is remembering

33 Identifying potential counterexamples of the latter is, of course, the original Wason's (1966) selection task. The improved performance for cheater-detecting is discussed in Cosmides and Tooby (1992); see Fodor (2000) for some criticism.

34 Cf. Kirchoff and Newsome (2012).

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

things. Since my purpose is only to illustrate a particular methodology for diagnosing REC, it will suffice to focus only on the possibility that each partner's remembering is radically extended into the other.) I will argue that epistemic evaluation points towards a verdict of "No." In so doing, I will illustrate the effectiveness of, at least sometimes, using the epistemic as a "mark" of the cognitive.

Before sketching a case for the "No" verdict in the next section, however, it is helpful to consider some of the various alternative interpretations of the case in question. To begin with, it might well be conceded that the remembering processes of the husband and wife are interestingly "coupled" *even though* the remembering processes of one do not radically extend into the other.[35] Consider that "coupling" might simply entail capacities for remembering are impoverished without the partner's aid. Presumably, remembering is a performance that each is capable of only in suitable circumstances. When memories are not particularly robust, these suitable circumstances for remembering might involve being in a state of mind typically brought about by *priming* or *prompting*. For instance, it might be that although the husband may have a stored memory that the pestering began on a Monday night, retrieval is typically only possible when the wife helps to "jog" his memory. If this is how we should understand the quoted case, then we need not think about it as a case of REC at all *even though* the husband's memory capacities are impoverished without the help of his wife (thus showing them to be "coupled"). The remembering might be happening entirely within the husband's head even though the background conditions required for this remembering process are typically brought about through conversation with the wife.

Another interpretation of the case worth considering is that at least some of the appearance of remembering in the case is, in fact, confabulation-like *fabrication*. For instance, it may be that insofar as the husband takes himself to be remembering, he is mistaken. Really, the husband believes on the basis of the wife's testimony, but perhaps this belief comes so naturally that the husband takes himself to be remembering. In fact, the husband has forgotten which day the pestering began, but gains a new belief about it upon hearing the wife.

The husband's remembering would be a *radically extended* cognitive process if the process is partly located in the wife. For instance, it could be that the husband's memory has been partly offloaded onto and subsequently stored by the wife. However, the practical necessity of the wife for the husband's remembering is not enough for REC. Both the priming and fabrication possibilities intimately involve the wife, and they have to be ruled out if the remembering process is to be established as radically extended. (I'm ignoring, of course, the further possibility that the case is one where the husband's memory is autonomous, i.e. he could simply remember everything on his own without the wife – and vice versa. Such a case does not appear to support the contention that the memories of the wife and husband are "coupled," much less that remembering is radically extended. So, for the sake of argument, I'm simply granting the case is not of this sort.)

6. EPISTEMIC EVALUATION AS A GUIDE TO COGNITION: A CASE STUDY

The important thing to notice is that it is *epistemically significant* whether the priming, fabrication, or radically extended cognition scenario obtains. For instance, if the husband

35 See Adams and Aizawa (2010) on the coupling fallacy.

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

is *fabricating*, there is a deficit in his self-*knowledge*, but if he is remembering (either through priming or radically extended cognition) there's no such (epistemic) deficit. More importantly for purposes here, where the remembering is happening has an impact on how we evaluate the resulting the belief. Suppose, for instance, that there is unusual activity located within the wife's head. If the husband's cognition radically extends into the wife, then this is potentially a problem with the husband's own cognitive processing that could be relevant to the quality of his cognitive performance, thus impacting the epistemic credentials of any resulting belief. If, however, the husband's cognition does not radically extend into the wife, then although he might be misinformed by the wife, his cognitive processing is itself unaffected.

To make these suggestions more vivid, imagine that the wife takes herself to be remembering the day of the cooking classes that she was taking when she met her now husband, but she confuses these cooking classes with wine tasting classes that she took in the same location a few months earlier. Indeed, she confuses them with the wine tasting classes in the sense that her memory of the day of the week of the cooking classes counterfactually depends on the day of the week of the wine tasting classes rather than on the day of the week of the cooking classes (as it should). Moreover, we might suppose she might even recognize this confusion by realizing that the classes she is remembering took place at a different time in the year than the classes she was taking when she met her husband. However, as it stands, she fails to notice the incoherence of her position. She simply reports that the day of the week of the cooking classes was Monday (because that was the day of the week of the wine tasting classes). In this scenario, the wife's belief token would be *rationally defective* in the sense that it is the product of a poor cognitive performance. She doesn't have an occurrent belief that she met her husband after her classes on Monday *because* this type of belief is evidentially supported. Indeed, this is true *even if*, in fact, she has evidential support for a belief that those cooking classes took place on Monday. For instance, it might be that if she thought about it more, she would realize that these cooking classes took place on the same day of the week as her wine tasting classes had. Still, her occurrent belief that the cooking classes took place on Mondays is not held (even partly) for this reason.

Arguably, if the husband's cognitive processes extended out into the wife's head, the problem we are imagining for the wife would straightforwardly be his as well; the husband's token belief would also be *rationally defective* in roughly the same sense: he would have this occurrent belief that the day was Monday, but not because he has evidential support for a belief of this type. (Shortly, I will defend this inference, but it is useful to consider the implications of it first.) On the other hand, if the husband's cognitive processes do *not* extend out into the wife, then this occurrent belief need not be rationally defective in the same way. If the husband has accepted the word of the wife as something he himself remembers – if he is *fabricating* its origin – then there is certainly something wrong with his second-order belief concerning the origin of his belief that the day was Monday. But, it isn't obvious that there is anything wrong with the first-order belief that it was Monday *itself*. This belief is (presumably) reasonably held because he is reasonably taking his wife's words on their face (even if he thinks he's remembering). Alternatively, if the wife's utterance simply "jogged" the husband's memory – so that he was *primed* to genuinely remember the day of the week – then the wife's confusion presents no problem for the husband at all. Her confusions are not his.

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

The upshot is that it appears that we can adjudicate the claim that the husband's remembering is radically extended by considering the epistemic status of his occurrent belief that the wife's classes were on Monday. If his occurrent belief inherits the epistemic status of the wife's, then there's radically extended remembering. If it doesn't, then there isn't. I contend that, in the kind of scenario envisioned where one partner is confused, this very confusion doesn't transfer to the other. For instance, if the wife's confusion comes to light, it is unnatural to attribute this confusion to the husband. This is especially so if the husband insists that he (now) remembers that it really was Monday and can even report supporting memories. However, it is still unnatural even if the husband defers to the wife and withdraws his statement that the classes took place on Monday once the initial confusion is revealed. Indeed, it seems much more natural to think that the husband has fallen prey to his gullibility (and perhaps lack of self-knowledge) than it is to think that he has fallen prey to confused remembering in the very same way as his wife has. The husband may well have exhibited an epistemic flaw, but it is a distinct one. (Certainly, the *husband* hasn't confused the day of the cooking classes with the day of the wine tasting classes even if the husband is, in fact, confused!) This epistemic assessment undermines the claim that the husband's remembering is radically extended into the wife – at least in a bad sort of case where there's a problem located in the wife's head. However, if there's no REC in the bad case, then there's no motivation for postulating REC in the good case where there's no problem at all. So, there's no apparent reason to think that there's REC in the original Sutton et al. (2010) case quoted in the last section.

This line of reasoning against REC can be summarized succinctly as an instance of *modus tollens*. From the premise (1) that *if the husband's process of remembering radically extends into the wife's head, he would be confusedly remembering if she were* and the premise (2) that *he would not be confusedly remembering if she were*, one infers that the husband's process of remembering is not radically extended into the wife's head. The obvious way to resist this line of reasoning is to deny the premise (2). When the wife confusedly remembers, it merely *seems* as if the husband isn't also confusedly remembering in roughly the same way. It might seem otherwise because the husband's belief isn't, in fact, rationally defective *even though* he is confusedly remembering because his remembering process is intertwined with hers. It might be argued, for instance, that the husband has two sources of evidential support: his remembering (which draws on resources located in the wife) and the wife's testimony (which doesn't). Moreover, he believes on the basis of evidential support from both sources at once. So, even if his believing by remembering is defective, his simultaneous believing by accepting testimony is not. In this way, the husband's belief is in better shape than the wife's (even though his remembering process is deeply parasitic on hers!).

In my view, this resistance should itself be resisted. If the wife were feeling especially suspicious, she could take an alienated perspective on her own memory faculties. She could view them as she would an external witness.[36] After all, these memory faculties largely operate in ways that aren't directly accessible from a first-person perspective just as external testifiers do. Thought of in this alienated way, the wife's memory faculties are "testifying" to her. Still, the wife has reason to think that they are generally reliable, and, on this basis, might accept the "testimony." In other words, it might be argued that

36 Cf. Ichikawa and Jarvis (2013) on self-experimentation.

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

the wife too has two sources of evidential support: her remembering and the "testimony." These sources are analogous to those of the husband. Moreover, she might believe on the basis of both sources at once – just as the husband does – so there should be no difference between these two beliefs. While this second source of evidential support might be a mitigating factor, it still seems that the wife's belief is defective. Ultimately, it is the product of confusion. If the husband's remembering process is radically extended into the wife, the same is arguably true of him.

A possible rejoinder is that testimony from the wife to the husband is genuine whereas there is only "testimony" – as the view from an alienated perspective – of the wife's memory faculties to her. It might be argued that this difference matters for the evaluation of their respective beliefs concerning the day of the cooking classes. However, this rejoinder is vexed for the advocate of REC. The difference between genuine testimony and "testimony" should only matter for epistemic evaluation if there's an important cognitive difference in the way these two inputs are received. However, if there is an important cognitive difference between the way that the husband treats the input from his wife and the way that the wife treats the input from her memory faculties, then there's a significant failure of "parity" – as discussed in Clark and Chalmers (1998).[37] The husband is no longer relying on the wife in the way that a subject relies on her memory faculties, so there's no apparent motivation for supposing that the wife is functioning as – and therefore is – a part of his memory system.

7. CONCLUDING THOUGHTS

The discussion in §6 is meant to be *suggestive* rather than *conclusive*. Perhaps, further examination will reveal the Sutton et al. (2010) case, in fact, to be a case of REC. My interest in this paper is not in arguing for or against REC, but rather in demonstrating a kind of *methodology* for doing so, a methodology that is grounded in the relationship between epistemology and REC discussed in the earlier sections of the paper. Because epistemological theory is framed in terms of the cognitive (as indicated in §2), but unaffected by the prospects for REC (as discussed in §§2–3), epistemological theories can be used as an effective standard for diagnosing REC. §§5–6 are simply intended to illustrate how this diagnosis might go. Our pre-theoretical tendencies towards epistemically assessing subjects in particular ways – e.g. tendencies not to think that confusions transfer between partners – may reveal pre-theoretical commitments about where the cognition of these subjects is located. Further systematic epistemic assessment may confirm these pre-theoretical commitments. Or, it may not. Either way epistemology provides a starting place for examining whether REC is present. This is the intended moral of this paper.[38]

37 See also Clark (2008: 77–8) and Wheeler (2010).

38 The moral, I think, might be generalized to apply even in cases where it is less obvious whether any particular belief is participating in the cognitive process in question. Consider, for instance, the example – from Clark and Chalmers (1998) and Clark (2008) – of Tetris. Notably, Tetris moves can made reasonably or not, depending on whether the cognitive process of selecting the move exhibits epistemic competency. Thus, we can evaluate whether the (external) rotating of a Tetris piece is part of the cognitive process of selecting a move by considering whether anomalies in the rotating change whether a move was made reasonably. For instance, if a move cannot be reasonably made merely by taking into account how the piece seems to rotate (due to hallucination), it might well indicate that cognition does extend to the rotating since the rotation is part of exercising the epistemic

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

Nevertheless, if there's a more specific lesson to be learned from thinking about the Sutton et al. (2010) case, it is this: while the prospect of recruiting external resources for cognition might provide wonderful opportunities for more extensive cognitive success, it simultaneously provides further locations for cognitive failures. Anywhere a subject's cognition is located is a place where the subject's cognition could be defective. Consequently, before judging that one's cognition extends to some particular place, it is worth considering whether one could be confused, muddleheaded, irrational, etc. in virtue of what happens there. My own view is that it is not generally plausible that any of us could be confused, muddleheaded, irrational, etc. in virtue of what happens where other human beings are located.[39]

REFERENCES

Adams, F. and Aizawa, K. 2010. *The Bounds of Cognition*. Malden, MA: Wiley-Blackwell.
Aizawa, K. 2012. 'Distinguishing Virtue Epistemology and Extended Cognition.' *Philosophical Explorations*, 15: 91–106.
Burge, T. 2007. *Philosophical Essays, Volume 2: Foundations of Mind*. Oxford: Oxford University Press.
—— 2010. *Origins of Objectivity*. Oxford: Oxford University Press.
Carter, J. A. 2013. 'Extended Cognition and Epistemic Luck.' *Synthese*, doi 10.1007/s11229-013-0267-3.
Carter, J. A. unpublished. 'Epistemic Luck and Propositional Memory.'
Carter, J. A. and Pritchard, D. Forthcoming. 'Extended Entitlement.' In P. Graham and N. Pedersen (eds), *Epistemic Entitlement*. Oxford: Oxford University Press.
Christensen, D. 2004. *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford: Oxford University Press.
Clark, A. 2008. *Supersizing the Mind*. Oxford: Oxford University Press.
Clark, A. and Chalmers, D. J. 1998. 'The Extended Mind.' *Analysis*, 58: 7–19.
Comesaña, J. 2006. 'A Well-Founded Solution to the Generality Problem.' *Philosophical Studies*, 129: 27–47.
Conee, E. 1992. 'The Truth Connection.' *Philosophy and Phenomenological Research*, 52: 657–99.
—— **and Feldman, R.** 1998. 'The Generality Problem for Reliabilism.' *Philosophical Studies*, 89: 1–29.
Cosmides, L. and Tooby, J. 1992. 'Cognitive Adaptations for Social Exchange.' In J. Barkow, L. Cosmides, and J. Tooby (eds), *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, pp. 163–228. Oxford: Oxford University Press.
Feldman, R. and Conee, E. 2001. 'Internalism Defended.' *American Philosophical Quarterly*, 38: 1–18.
Fodor, J. A. 2000. *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology*. Cambridge, MA: MIT Press.
Foley, R. 1993. 'What I am to Believe?' In S. Wagner and R. Warner (eds), *Naturalism: A Critical Appraisal*. Notre Dame, IN: Notre Dame University Press.
Goldman, A. 1976. 'Discrimination and Perceptual Knowledge.' *Journal of Philosophy*, 73: 771–91.
—— 1979. 'What Is Justified Belief?' In G. Pappas (ed.), *Justification and Knowledge*. Dordrecht: Reidel.
Greco, J. 2003. 'Knowledge as Credit for True Belief.' In M. DePaul and L. Zagzebski (eds), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, pp. 111–34. Oxford: Oxford University Press.

competency. (For an argument in this vein, see Jarvis (2014) on pen and paper in arithmetic.) Thanks to an anonymous referee for raising this point.

39 Thanks to J. Adam Carter and an anonymous referee for their thoughts and suggestions.

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press

—— 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.
Hare, R. M. 1952. *The Language of Morals*. Oxford: Clarendon Press.
Horwich, P. 2006. 'The Value of Truth.' *Noûs*, 40: 347–60.
Ichikawa, J. J. and Jarvis, B. 2013. *The Rules of Thought*. Oxford: Oxford University Press.
Jarvis, B. 2013. 'Knowledge, Cognitive Achievement, and Environmental Luck.' *Pacific Philosophical Quarterly*, 94: 529–41.
—— 2014. 'Evaluating the Extended Mind.' *Philosophical Issues*, 24: 208–28.
Kelp, C. 2013. 'Extended Cognition and Robust Virtue Epistemology.' *Erkenntnis*, 78: 245–52.
Kirchoff, M. D. and Newsome, W. 2012. 'Distributed Cognitive Agency in Virtue Epistemology.' *Philosophical Explorations*, 15: 165–80.
Koksvik, O. Unpublished. 'Intuition, Belief, and Rational Criticisability.'
Lynch, M. 2004. 'Minimalism and the Value of Truth.' *Philosophical Quarterly*, 54: 497–517.
—— 2009. *Truth as One and Many*. Oxford: Oxford University Press.
Nozick, R. 1974. *Anarchy, State, and Utopia*. New York, NY: Basic Books.
Pritchard, D. 2005. *Epistemic Luck*. Oxford: Oxford University Press.
—— 2010. 'Cognitive Ability and the Extended Cognition Thesis', *Synthese*, 175: 133–51.
Putnam, H. 1975. 'The Meaning of "Meaning".' *Minnesota Studies in the Philosophy of Science*, 7: 131–93.
Quine, W. V. O. 1953. 'Two Dogmas of Empiricism.' In *From a Logical Point of View*. New York, NY: Harper and Row.
Rowlands, M. 1999. *The Body in Mind*. Cambridge: Cambridge University Press.
Rupert, R. D. 2009. *Cognitive Systems and the Extended Mind*. Oxford: Oxford University Press.
Shapiro, L. 2011. *Embodied Cognition*. New York, NY: Routledge.
Sosa, E. 2003. 'The Place of Truth in Epistemology.' In M. DePaul and L. Zagzebski (eds), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, pp. 155–80. Oxford: Oxford University Press.
—— 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Vol. 1*. Oxford: Oxford University Press.
Sosa, E., Kim, J., Fantl, J., and McGrath, M. (eds). 2008. *Epistemology: An Anthology*. Oxford: Blackwell Publishing.
Sutton, J., Harris, C. B., Keil, P. G., and Barnier, A. J. 2010. 'The Psychology of Memory, Extended Cognition, and Socially Distributed Remembering.' *Phenomenology and the Cognitive Sciences*, 9: 521–60.
Vaesen, K. 2014. 'Dewey on Extended Cognition and Epistemology.' *Philosophical Issues*, 24: 426–38.
Wason, P. C. (1966). 'Reasoning.' In B. Foss (ed.), *New Horizons in Psychology*, pp. 135–51. Harmondsworth: Penguin.
Wilson, R. 1994. 'Wide Computationalism.' *Mind*, 103: 351–72.
Wheeler, M. 2010. 'In Defense of Extended Functionalism.' In R. Menary (ed.), *The Extended Mind*, pp. 245–70. Cambridge, MA: MIT Press.
Whiting, D. 2012. 'Does Belief Aim (Only) at the Truth?' *Pacific Philosophical Quarterly*, 93: 279–300.
—— 2013. 'The Good and the True (or the Bad and the False).' *Philosophy*, 88: 219–42.
Williamson, T. 2002. *Knowledge and its Limits*. Oxford: Oxford University Press.
Zagzebski, L. 1994. 'The Inescapability of Gettier Problems.' *Philosophical Quarterly*, 44: 65–73.
Zangwill, N. 2005. 'The Normativity of the Mental.' *Philosophical Explorations*, 8: 1–19.
—— 2010. 'Normativity and Metaphysics of Mind.' *Australasian Journal of Philosophy* 88: 21–39.

BENJAMIN JARVIS received his PhD in philosophy from Brown University in 2010. He has published on the epistemology of philosophy, knowledge, epistemic value, virtue epistemology, mental content, and the extended mind. He is currently studying the role of data analysis in practical decision-making – particularly under risk – at the New York University, Stern School of Business.

https://doi.org/10.1017/epi.2015.30 Published online by Cambridge University Press