Patrolling the mind's boundaries

Abstract: Defenders of the extended mind thesis say that it is possible that some of our mental states may be constituted, in part, by states of the extrabodily environment. Often they also add that such extended mentation is a commonplace phenomenon. I argue that extended mentation, while not impossible, is either nonexistent or far from widespread. Genuine beliefs as they occur in normal biologically embodied systems are informationally integrated with each other, and sensitive to changes in the person's overall system of beliefs. Environmental states, however, fail to satisfy this central feature of the functional role of belief, and hence fail to be genuine mental states.

1. Introduction

Traditional materialism claims that mental states are located in the brain and body, if anywhere. In recent years, however, some materialists have argued for an apparently radical claim about where we should draw the mind-world boundary. They suggest that, in many cases, the boundary should be pushed outward to include parts of the extrabodily environment within the mind itself. As John Haugeland puts the point, "intelligence itself abides 'out' in the world, not just 'inside'—contra cognitive science, classical or otherwise" (1998, p. 232). According to this *extended mind thesis,* at least some mental states may be partially or wholly constituted by states outside of the brain and body. Laptop computers, notebooks, and maps—not to mention calculators, slide rules (Hutchins, 1995), the road signs by which we steer, even natural language itself (Dennett, 1996)—may literally realize parts of our minds.

My concern here will be primarily with the claim that environmental states might constitute mental states, in something like the commonsense conception. Addressing this claim requires some notion of what the commonsense conception of mental states is actually committed to. I will argue that if we attend to the functional and explanatory roles this conception describes, it appears unlikely that our minds are, in fact, extended into the extrabodily environment. The extended mind thesis makes a modal claim: it is *possible* for mental states to

be externally constituted. Arguing its falsity would require arguing that it is necessary that beliefs are internal states of the organism that possesses them. I don't have such an argument. For all I will say here, the thesis might be true. But even if the extended mind thesis is true, the conditions on minds being extended are far more stringent than have been advertised. And this conclusion is of independent interest, since defenders of extended minds often think that extended mentation is a virtually commonplace phenomenon, particularly among contemporary, literate, technology-using people.[1]

2. The case for extended minds

The extended mind thesis commits us to the claim that "*beliefs* can be constituted partly by features of the environment, when those features play the right sort of role in driving cognitive processes" (Clark & Chalmers, 1998, p. 12). This applies just as well to other mental states, but for simplicity I will mostly focus here on belief. The argument for this thesis rests on a simple, orthodox functionalist principle. The principle is that "[w]hat makes some information count as a belief is the role it plays, and there is no reason why the relevant role can be played only from inside the body" (Clark & Chalmers, 1998, p. 14).[2] This principle (or at least its first conjunct) is no more than the elementary functionalist tenet that to be a belief is just to play the causal/functional role of belief.

Human and animal brains, and perhaps the odd computer, are assumed to be the normal realizers for mental states understood functionally. It has long been noticed that functionalism

[1] Andy Clark, for instance, says: "This kind of broadening is probably most plausible in cases involving the external props of written text and spoken words, for interactions with these external media are ubiquitous (in educated modern cultures), reliable, and developmentally basic …. Much of what we commonly identify as our mental capacities may … turn out to be properties of the wider, environmentally extended systems of which human brains are just one (important) part" (1997, p. 214).

[2] I'd note that the principle requires slight revision, since it is not information *per se* that counts as a belief, but a state that carries or represents that information. Beliefs are not information itself, they are bearers of information.

allows peculiar systems whose states realize the role of belief: Martian hydraulic brains, the nation of China, a well-organized beehive, etc. Neural prosthetics such as the silicon-based visual cortex (Searle, 1992) provide examples of systems containing hybrid realizers (both neural and silicon), and a series of intermediate cyborg cases connects the purely meat brain with the purely metal brain. Unusual realizers are a staple of the functionalist literature. The hybrids described by advocates of extended minds differ only in lying outside of the normal brain-body system. This is an admittedly novel possibility that has not hitherto received much attention in the literature. But in itself it is simply an instance of an unusual realization of a mental state, and thus a fixture familiar to functionalists.

So there is a sense in which the extended mind thesis should not be seen as especially radical. Functionalists have all along been committed to the possibility of extrabodily states playing the role of beliefs and desires. But is the extended mind thesis true of us? Is it true that some, perhaps many, of our mental states are constituted (wholly or in part) by what are traditionally regarded as environmental states?

The answer depends on whether there are external states that satisfy the functional role of those mental states. Clark and Chalmers develop one vivid case that allegedly illustrates such a situation: the case of Otto. Otto suffers from Alzheimer's disease. He cannot encode, retain, or recall new information, and much of the information he previously retained and could recall is now lost to him. But he does have a notebook that he carries with him constantly, in which he writes down new information he thinks will be useful to him later. He is in continuous touch with the notebook and its contents, and he relies on it just as those with normal (biologically based) memories rely on them to navigate the physical and social environment. If the notebook were taken from him, Otto would suffer; he would be unable to carry out many of the acts he now can

carry out, such as getting to the store successfully. His access to the information in the notebook is reliable, even if not as apparently unmediated as our access to our own beliefs and memories. And he endorses and acts on the information in the notebook when he retrieves it, just as we normal subjects use our beliefs to guide our behavior. Hence, the notebook and its contents fulfill the functional role of belief for Otto, and by the general functionalist dictum, some of his beliefs must be (partially) constituted by states of the environment.

The conclusion that the notebook partially constitutes some of Otto's beliefs turns crucially on the claim that it (and its contents) fulfill the functional role of belief. Clark and Chalmers lay out four conditions that the notebook satisfies that align it with beliefs as we ordinarily regard them:

> First, the notebook is a constant in Otto's life …. Second, the information in the notebook is directly available without difficulty. Third, upon retrieving information from the notebook he automatically endorses it. Fourth, the information in the notebook has been consciously endorsed in the past, and indeed is there as a consequence of this endorsement. (1998, p. 17)

I don't dispute that Otto's notebook fulfils the functional role sketched here. I think, however, that the notebook lacks a key feature that is part of the normal functional role of belief, and that is absent from the above list of features. If this is true, and if the argument for the extended mind thesis turns on the functionalist principle, then the case of Otto is not in fact a case of extended mentation.

3. Against extended minds: The role of belief

The feature of belief that I will focus on is the fact that particular beliefs are sensitive to changes that occur elsewhere in a person's mental economy. Beliefs are, as I will say, normally *informationally integrated* with, and updated in concert with, other beliefs (and further mental states of the subject, such as desires). But most alleged cases of externally located mental states do not share this feature. So, by the functionalist principle, they cannot be beliefs.

Informational integration is a relation that sets of mental states bear to one another such that a change to one such set causes appropriate changes in other, relevantly related sets of states. To count as a belief a state has to be part of a system of states in which processes of integration and updating function to keep the subject's mental contents in epistemic equilibrium to some degree or other. The hard work of spelling out how informational integration works requires saying under what circumstances a set of states is *relevantly* related to another and what changes count as *appropriate*. Inasmuch as this task relies on a substantially fleshed out empirical theory of human reasoning, it is far beyond the scope of this paper. Here I will just say a bit about how informational integration operates in quotidian cases.

Take two sets of mental states: my beliefs about people's marital status, and my beliefs about people's home addresses. There are some obvious relations between these two sets of beliefs. For one thing, if I believe that two people are married, I will (in the default case) believe that they live at the same address. This default can be overridden, but it is not a bad *prima facie* assumption. So if I believe that Sam and Max are married, I'll typically believe they share an address. If my belief about their marital status changes, some of my beliefs about their addresses will probably change, too. At least one of them will probably not have the same address any more—it depends on who gets to keep the rent-controlled apartment. This is a banal illustration

of informational integration, but its banality just emphasizes the fact that integration is part of the everyday dynamics of belief.[3]

Informational integration is normally achieved by mechanisms that operate below the threshold of our conscious awareness and control. Moreover, these mechanisms typically operate rapidly, allowing many beliefs to be updated in the same relatively brief span of time. But neither of these facts about integration are *necessary* to the process itself. We may sometimes update aspects of our belief store ourselves, reasoning through what needs to be changed in a conscious, effortful, and comparatively slow fashion. When I speak of informational integration as being rapid and effortless, then, I don't mean to imply that it is never otherwise, nor that this is the essence of the process.

Two situations illustrate how the property of informational integration fails to be respected by so-called extended minds. In a *mixed* external memory situation, a subject has two relevantly related beliefs A, stored internally, and belief B, stored externally. In a *pure* external memory situation, a subject has two beliefs, A and B, both stored externally. Each situation demonstrates how externally stored memories can fail to satisfy the functional role of belief.

We can exemplify a mixed external memory situation with the case of Waldo, who suffers from memory loss, but of a less severe kind than Otto's. His is intermittent, but he is aware of it, and he carries a notebook in order to record information he thinks he will later need to remember. But he does not write down everything. Some things he tries to remember himself,

[3] Informational integration has been extensively discussed by Fodor (1983; 2000). In the closing sections of (1983), Fodor suggests that central cognitive systems—those non-modular parts of the architecture containing the organism's beliefs and desires—have the property of *isotropy*, meaning (roughly) that, in principle, a doxastic state may gain or lose support depending on states taken from anywhere in the totality of the organism's other doxastic states. As Fodor colorfully puts it (1983, p. 107), our botanical beliefs are in principle constrained by our astronomical beliefs. In my terms, all that keeps these two stores of information from being integrated is the fact that we aren't (at least, I'm not) presently aware of the linkages that would make one body of information relevant to the other.

and sometimes he succeeds in this. Now suppose that the address of the museum is one thing he thinks he may forget. Hence he writes in his notebook the sentence "The museum is on 53rd Street". He learns later that the museum has been torn down to make way for a bypass, and this fact he remembers biologically.

I take it that the normal subject who believes at t_1 that the museum is on 53rd Street will, on learning at t_2 that the museum has been torn down, no longer believe that. She may believe instead that the museum *was* on 53rd Street, that the museum is no more, and so on. Acquiring relevant new information causes changes in other related beliefs in the agent's mental economy. These changes form part of the normal causal role of belief. Beliefs tend to be updated to reflect novel information the believer acquires.

But notice that some of Waldo's beliefs are not updated in this way. In particular, the beliefs that are externally stored are not automatically and unconsciously updated to reflect his new information about the museum. Presumably, if Waldo also has stored in his head the belief that the museum café makes a good latte, then when he learns that the museum has been demolished, he will no longer believe that (or will change it to the past tense). The notebook sentence "The museum café makes a good latte" will not similarly be updated or altered as a reliable consequence of the state of Waldo's head having changed, however.

This example illustrates the point that, normally, a person's beliefs are stored as part of a system that actively propagates changes to the other contents of the system. Making systems that do this in a sensible, relevant, timely fashion is a major challenge for artificial intelligence; it is arguably at the heart of the frame problem, for instance. But if being subject to this sort of updating is part of the functional role of belief, then the notebook states cannot be beliefs of Waldo's, since they are not updated.

I should add that the beliefs under discussion here are best described as a person's standing beliefs, not her occurrent beliefs. They are the contents of some portion of her long term memory store. Furthermore, all of the examples here are to be construed as cases of *explicit* belief. There are complex issues concerning the legitimacy of the concept of an implicit or tacit belief that take us too far from the concerns raised here (see Lycan, 1981). Employing a crude simplification, we can think of explicit beliefs of both the internal and the putatively external sort as being sentences in some language—whether it be English or Mentalese—that play a certain functional role. The correctness of this linguiform model of belief is not at issue here. Skeptics about the model should replace sentence-like structures with any representational vehicles that appeal more to them. The arguments go through just as well, so long as the representational system chosen permits explicit representation.

Consider now a pure external memory situation, in which both pieces of information are stored in the world. Otto may have written on page 10 of his notebook that the museum is located on 53rd Street, and on page 20 that the museum has been demolished to make way for a bypass. Do these states count as Otto's beliefs? Arguably not, for there is no reason to think that it is a normal, automatic, consequence of writing the second sentence that the first sentence will be updated or erased to reflect his new total informational state. The informational inertness of this sort of 'external memory' is precisely what should disqualify it from being a literal belief store. Interestingly, Clark elsewhere criticizes Herbert Simon for adhering to an "overly passive (mere storage) view of biological memory" (Clark, 2001, p. 140). My claim is that the extended memories so far examined are all far too inert to contain beliefs.

Notice that if we suppose these external states to constitute Otto's standing beliefs, then we must say that he both believes that the museum is on 53rd and that it has been demolished.

Minds that are extended in Otto's sort of way seem especially prone to falling into this sort of inconsistency and irrationality. The source of such error is the absence of the normal consistency-maintaining doxastic mechanisms located in the functional architecture of believers. Either, then, the external states are not beliefs, or they are the beliefs of a particularly irrational agent. Otto and Waldo's belief stores contain obviously inconsistent propositions. And, moreover, these inconsistencies do not require a massive chain of deductions or refined logical acumen to detect. To overlook their inconsistency if they were placed side by side would require a bout of serious inattention.

Why not opt for the latter of these options and plead irrationality for the extended thinkers? One might argue that if a subject is too irrational, it becomes impossible to attribute beliefs sensibly to her at all; perhaps beliefs can only exist in substantially rational agents (Dennett, 1987). I will not take this tack here. Rather, I would point out that if extended minds are extremely common, perhaps ubiquitous, then this sort of irrationality is also likely to be extremely common. For it is extremely easy to be 'coupled' to an information source that contains such inconsistencies. All that it takes for Waldo to be in this situation is for him to have forgotten that he previously entered something relevant to the location of the museum in his notebook. Otto is in a still worse position, given that he is even more reliant on his notebook.

The larger the external source of information, the less likely a person is to recall that related information has previously been entered. On a single page of note paper, I am unlikely to miss obviously related information. In an entire notebook, let alone the files of my laptop computer, I am much more likely to do so. Of course, increased size also makes such external stores more useful. So there is a grim trade-off here: if I am regarded as being cognitively coupled to an entire store of information whenever *part* of it fulfils the functional role of a belief

state for me, then as these stores become larger and potentially more useful to me, I run a progressively greater risk of falling into doxastic incoherence. It is difficult to see what sort of explanatory gain we might make by positing that there is widespread irrationality of this kind. Since advocates of extended minds argue that they *will* bring some such gains to our everyday and scientific explanatory practices, they are unlikely to embrace this option.

Can one imagine a situation such as Waldo's in which the notebook's contents are updated as the subject acquires new information? Certainly. Perhaps when Waldo writes a new sentence in his notebook he checks each page to see whether it needs to be changed in light of what he is now writing. Perhaps he regularly flips through the pages checking whether what is written there conforms with what he has stored in his head. He might even implement a regular system of indexing and updating information in this way. This would be a conscious, effortful way of doing what the cognitive system frequently does unconsciously and rapidly. Would this satisfy the informational integration requirement?

I presume that whatever algorithm the cognitive system is designed to use in carrying out its updates can be written down and explicitly used by an agent to update the contents of his notebook and bring it in line with the rest of his beliefs. (Although I'd add that we have hardly any substantial theory of what this algorithm might be.) Alternatively, a person might devise a different system of updating that has the same effects. It matters that updates happen, not that they happen by one particular kind of mechanism or other. I have already said that sometimes we update our intracranially realized beliefs in a self-conscious, effortful fashion. So it's not impossible for external states to be integrated in the way other beliefs are. Since my concern here is only to show that *de facto* it is massively unlikely that there are any extended minds, this bare metaphysical possibility is no threat.

Notice, however, that if this is the possibility the defenders of extended minds embrace, it raises the bar *considerably* on what it takes for an external system to contain states that count as a subject's beliefs. For the requirement now involves a person's sorting through a mass of external media, tracking down relevant connections, and making appropriate adjustments consciously and explicitly. Innovations like indexes may reduce the effort, but indexes introduce further burdens of their own, such as compiling and updating the index itself. But the main point is that subjects who carry out such elaborate updates are patently far less common than those who merely store information in the environment. Indeed, perhaps there are no such subjects. Otto himself, as originally described, does not qualify. Thus, when we focus on what sort of work it would take to consciously and explicitly adopt a strategy of updating external information stores, extended minds appear far less commonplace than the homely case of Otto and his notebook might initially suggest.

This isn't to say that extended minds might not someday be realized. Consider the case of Wanda, who is somewhat more high-tech than her friend Waldo. Wanda does not use a notebook, but instead has a portable electronic memory device that is directly jacked into her brain. The device is seamlessly integrated with her natural cognitive architecture; it implements exactly the same basic cognitive processes and stores information in the same representational format as her biological memory does. Moreover, the information encoded in the device is integrated, via the continuously active link, with her biological memories. Nevertheless it is located outside of the brain-body system as it is usually conceived. Such devices are certainly impossible given current technology, and perhaps always will be. But who knows? The point here is just that although such devices are not available to us, their states would be far better candidates for externally stored beliefs than the states of Otto's notebook, since they are *actively*

functionally integrated with the rest of Wanda's belief system. No environmental states that we interact with satisfy this condition.

We should be careful about interpreting the demand that beliefs be subject to informational integration, however. We can clarify the proposal by contrasting it with a challenge to advocates of extended *cognition*, due to Adams and Aizawa (2001). They argue against the hypothesis that cognitive processing extends into the environment, on the grounds that different mechanisms are at work when we interact with the environment than when we use our natural, biological cognitive resources. For example, doing mental arithmetic uses one set of representations, processes, and capacities, while doing arithmetic on paper uses a quite different set. Cognitive processes, the proper subject matter for a science of psychology, are individuated in terms of underlying causal mechanisms. Where these mechanisms are different, as they are in the case of mental versus pencil-and-paper arithmetic, we cannot speak of the same process being carried out in two different systems. Given these differences in underlying causal mechanisms, they conclude, we should not think of human cognition as being extended into the environment.

Unlike Adams and Aizawa, I am not inclined to require that external states be updated using processes and mechanisms that are especially similar to those used by our native, biological belief-regulating systems.[4] I think that we *might* usefully regard some external states as beliefs even if the means by which they are updated differ from the means by which our own states are updated. The functional role of belief that I have been sketching here demands that beliefs be updated with the impact of relevant (or perceived to be relevant) information, in a

[4] Again, I note that Adams and Aizawa are arguing against the claim that cognitive processing might be extended into the environment, not the extended mind thesis as I have construed it. I am here entertaining the possibility that someone might apply a form of their argument to the extended mind thesis as well. Such an extension would involve the claim that extended minds are unlikely, since the causal processes that update beliefs in the head are unlikely to be duplicated outside of the head.

fashion that conforms to the expectations embodied roughly in commonsense psychological attribution and prediction. But this role is itself neutral among the many ways that these processes of updating might be realized in a computational and physical system. Our cognitive architecture uses a particular sort of algorithm to do its belief maintenance, but not everything that has beliefs must use precisely that algorithm. It is a matter of *implementation* which algorithm a particular individual or species employs.

This is just to say that beliefs and other folk mental states are not specified in terms of the precise computational mechanisms that regulate their interactions. At least, they are not so specified as far as commonsense functionalist descriptions go: folk mental concepts probably do not go that deep. The regularities that obtain among mental states specified in folk terms might be implemented in cognitive systems having different designs. As long as these higher-level regularities are preserved, we have a *prima facie* case for saying that a system has beliefs.[5] So whether a subject uses the same algorithm in updating her externally stored memories as she (or her mind/brain) does in updating her internally stored memories is not relevant to whether the external states are beliefs, although it might be relevant to whether the same sort of cognitive processing is going on inside and outside the head.

One might argue that informational integration is not a necessary feature of belief. For one thing, there are stubborn beliefs that no evidence seems likely to dislodge. People persevere in asserting all sorts of things in the face of apparent counterevidence. Do they therefore not really believe these things? These cases are not counterexamples to my claim, for the claim is

[5] I should say that the case is only *prima facie*, since folk psychology does not appear to be entirely neutral as far as questions of implementation are concerned. I take this to be the lesson of many people's intuitions that so-called 'Blockheads' (systems that behave identically to you or I, but whose behavior is entirely scripted by a gargantuan look-up table; see Block, 1995) lack beliefs and other mental states. If the internal architecture of a system lacks a requisite amount of complexity, it isn't typically counted as a believer; but it would be surprising if commonsense psychology had an opinion as to whether belief-change processes must be handled by one or another specific kind of algorithm. It is easy to exhaust our commitments to particular details of the implementation of our mental lives.

that a subject's beliefs are, in the typical case, open to automatic, largely effortless revision and amendment in the face of acquisition of novel information that the subject herself takes to be relevant to those prior beliefs. If a subject stubbornly, perhaps even irrationally, thinks that there is no evidence that would disconfirm one or another belief of hers, she nevertheless has that belief. (Fanatics are not called 'true believers' for nothing.)

Even if we grant this point, though, should we accept the claim that informational integration is essential to the functional role of belief? Could it not be a merely contingent feature of belief, perhaps just a typical accompaniment?

I think that it is independently plausible that informational integration is part of the more central portion of the role of belief. But functionalists have not directed much effort towards making the functional roles of particular mental states explicit, beyond mention of platitudes concerning pain and the so-called 'belief-desire law'.[6] Believers in extended minds appear to be committed to a kind of revisionist commonsense functionalism. They are willing to accept some modifications of our normal or typical characterization of beliefs, so long as other aspects are preserved and the overall explanatory role of belief remains unaffected. They might simply claim that, whether informational integration is in fact central or not, it is dispensable for the explanatory purposes that are served by positing beliefs. If so, then it is not a strike against external states that they are not integrated and updated as intracranial belief states are.

The trouble with this approach is that it does not seem to be true that informational integration is dispensable from the point of view of belief's explanatory function. One thing that our ability to attribute beliefs lets us do is to predict the changes that new information will bring about in someone's mental economy. If I tell you that the museum has been demolished, you no

[6] The 'belief-desire law' states, roughly, that if S desires P and believes that doing Q will conduce to getting P, then all things being equal S will do Q when she believes herself able to do so. Making this precise and contentful is clearly a major challenge.

longer believe that it is on 53rd Street. If I tell you that I have dyed my hair purple, you will cease to believe that its color matches my eyes. If I tell you that I have moved to a different state, you no longer automatically expect that I have the same phone number—yet the state of your address book does not change. Any number of our expectations about how subjects will respond to new information, and will change their own plans and expectations, depend on beliefs retaining their normal dynamics. So it isn't clear that we can detach the sensitivity of beliefs to other beliefs and retain the state's explanatory utility. A great deal of this utility is bound up with our ability to take the existence of changes such as these for granted.

This objection also recalls the earlier point that extended minds are liable to fall into certain kinds of irrational or incoherent belief. The point here is in fact stronger: given the lack of reliable causal processes that implement updating of external states in the appropriate way, so-called extended minds *default* to having potentially incoherent systems of belief. If we must literally attribute to extended minds these irrational belief systems, what explanatory power do we gain? What consequences for behavior or for future mental states can we derive from knowing that someone both believes that the museum is on 53rd Street and that the museum has been demolished? It just isn't clear *how* we might expect such a person to act, or what inferences we might expect him to draw. So from the point of view of folk psychology, broadening our concept of belief in a way that allows extended minds, many of which may be *de facto* irrational, comes with a price in terms of explanatory power. This point is especially serious for extended mind theorists who self-consciously advocate the thesis as part of a normative program of overturning our everyday conception of our minds and selves.

Still, this point might be pressed somewhat further. One might argue that everyday thinkers are in fact far less epistemically well-integrated than the informational integration

requirement demands. If we adopt too stringent a standard of integration, we run the risk that everyday believers won't make it over the bar either. In formulating this objection it may be worth considering briefly mental states besides belief. For example, the functional role of desire includes the property of representing the world as the agent would like it to be: as, for example, a world in which the agent has a cool milkshake. But desires, too, are states that are typically integrated both with other desires and with other kinds of mental states.

Consider how desires are integrated with beliefs. An obvious example of this relationship is that if I come to believe that P, I no longer desire that P be the case. The belief that the state of affairs desired has been satisfied brings the desire to an end. We can easily see how this fact about desire-belief integration could lead to trouble for extended minds. Recall the mixed internal/external case considered earlier; an analogous case can be constructed using a belief/desire pair. Suppose that one has an internally realized desire D and an externally realized belief B, and suppose that B is the belief that the state of affairs desired in D is satisfied.[7] If B is stored in some extracranial system to which the agent happens to be coupled (a notebook, laptop, etc.), then it will likely not be the case that the presence of the belief B will cause the desire D to cease. Unless there is a reliable causal mechanism—an updating process—bringing these states into their normal, integrated relationship, it seems that most every time an extended agent has a

[7] In the example, I have chosen to make the desire internal because it seems, for me at least, rather hard to conceive of desires as being externally stored states, although the full-blown version of the extended mind thesis does allow that it is possible. In part this difficulty seems to stem from the fact that the relative causal isolation of extracranial states makes them poor candidates to fill the motivational role of desires. Desires seem to need to have their 'hands' closer to the causal levers of action. In part the difficulty may also owe to the fact that desires have a phenomenological component—roughly, the component that consists of the felt *need* or *urge* for the thing desired—that is absent in the case of belief. Clark and Chalmers wish to restrict extended mental states to the non-phenomenological, so it is concessive to set the example up in this way. On the other hand, if one doesn't have such worries about the possibility of extracranially motivating states or about the phenomenology of desire, the example can just as easily be run with the desire being external and the belief being internal. The fundamental point about integration holds either way.

belief-desire pair that straddles the body-world boundary in this way, the pair will fail to respect the normal functional roles of both belief and desire.

The point of considering the example of desire is, first of all, to show that informational integration is not necessarily restricted to beliefs, although the form that integration takes is sensitive to the overall role of the kind of mental state being considered. More importantly, it is to drive home the fact that integration is not necessarily a demanding property, but one that is illustrated in countless everyday instances of mental state dynamics. It is the fact that purportedly extended mental states *reliably and by default* violate this minimal integration standard that renders them poor candidates to fill the functional and explanatory role delimited by our everyday concepts of the mental.

4. Conclusion

Beliefs and other mental states are embedded in a system that somehow brings it about that these states are updated and integrated with one another. They are not just acquired and inertly stored, waiting to be retrieved; even in storage, they are subject to processes that modify and sometimes eliminate them. Obviously, we are far from perfectly consistent. It is no part of my claim that these belief-modification processes lead us to be anything like ideally rational systems. But it is nevertheless true that new information tends to alter relevant related beliefs in appropriate ways. External memories such as Otto's notebook don't satisfy this condition by themselves, as Waldo's case makes clear. So such commonplace external memories are not cases of externalized beliefs. And much the same could be said of other kinds of mental states, such as desires, that have received less attention from advocates of extended minds. This raises the bar

that putative cases of extended mentation must clear. For now, for most of us, the mind stays safely within the boundaries of the body and brain.

Acknowledgements

I would like to thank Fred Adams, Ken Aizawa, Andy Clark, Brie Gertler, Rob Wilson, and Eric Winsberg for helpful discussion and comments on many previous drafts of this paper. Thanks also to the referees for this journal for their very thoughtful comments and advice.

References

Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology, 14,* 43-64.

Block, N. (1995). The mind is the software of the brain. In E. E. Smith & D. Osherson (Eds.), *An Invitation to Cognitive Science* (2nd ed., Vol. 3, pp. 377-425). Cambridge, MA: MIT Press.

Clark, A. (1997). *Being There*. Cambridge, MA: MIT Press.

Clark, A. (2001). Reasons, robots, and the extended mind. *Mind and Language, 16,* 121-145.

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis, 58,* 7-19.

Dennett, D. C. (1987). True believers. In *The Intentional Stance* (pp. 13-35). Cambridge, MA: MIT Press.

Dennett, D. C. (1996). *Kinds of Minds*. New York: Basic Books.

Fodor, J. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.

Fodor, J. (2000). *The Mind Doesn't Work That Way*. Cambridge, MA: MIT Press.

Haugeland, J. (1998). Mind embodied and embedded. In *Having Thought* (pp. 207-237). Cambridge, MA: Harvard University Press.

Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press.

Lycan, W. (1981). Tacit belief. In R. Bogdan (Ed.), *Belief: Form, Content, and Function* (pp. 61-82). Oxford, UK: Oxford University Press.

Searle, J. (1992). *The Rediscovery of the Mind*. Cambridge, MA: MIT Press.