The Extended Mind Thesis and the Causal Role of Beliefs: A Response to Weiskopf

Miriam Kyselo and Sven Walter

Institute of Cognitive Science

University of Osnabrueck

Albrechtstraße 28

D-49069 Osnabrueck, GERMANY

s.walter@philosophy-online.de

"The great end of life is not knowledge but action" (Huxley 1895, p. 422)

About a decade ago, Andy Clark and David Chalmers launched a forceful attack on 'brainbound' approaches to cognition which treat the mind as an essentially inner, i.e., neurally realized, phenomenon (Clark and Chalmers 1998). The traditional model of the mind as the brain, they held, is defective and must be replaced by one according to which our minds 'extend' into the environment. Our minds need not be located wholly within our bodies' boundaries in the sense that possibly "the human organism is linked with an external entity [like a sheet of paper and a pen, a pocket calculator, a notebook etc.; M.K. & S.W.] in a two-way interaction, creating a coupled system that can be seen as a cognitive system in its own right" (Clark and Chalmers 1998, p. 2)[i]

Clark (2008) describes this extended mind thesis as follows:

> [T]hinking and cognizing … may (at times) depend directly and non-instrumentally upon the ongoing work of the body and/or the extra-organismic environment. … [T]he actual local operations that realize certain forms of human cognizing include inextricable tangles of feedback, feed-forward and feed-around loops: loops that promiscuously criss-cross the boundaries of brain, body and world. The local mechanisms of mind, if this is correct, are not all in the head. Cognition leaks out into body and world. (Clark 2008, p. 000)

The suggestion that higher cognitive processes like, e.g., belief, memory, or learning may be extended has sparked a controversial debate. Daniel Weiskopf (2008) has lately attacked the extended mind thesis with regard to beliefs (for other lines of criticism see, e.g., Adams and Aizawa 2001, 2008; Rupert 2004). He argues that, since alleged cases of 'extended beliefs' lack a characteristic feature of beliefs properly so called (newly acquired beliefs are usually integrated with already existing beliefs rapidly, automatically and unconsciously), they do not count as genuine beliefs. As far as beliefs are concerned, he concludes, "the mind stays safely within the boundaries of the body and brain" (Weiskopf 2008, p. 275).

We believe that Weiskopf is wrong. Beliefs properly so called must not usually be integrated with a believer's system of already existing beliefs rapidly, automatically, and unconsciously. After some preliminary stage setting in section 1, section 2 suggests an alternative account of informational integration which is compatible with externally stored beliefs' being beliefs properly so called, emphasizes the crucial role that action

plays for cognition and thus stresses the embodied and situationally embedded nature of human cognizers as dynamical systems whose cognitive states and processes rely on, and are shaped by, their interaction with the environment. Section 3 argues that even if informational integration is de facto usually rapid, automatic, and unconscious, this is not an essential feature of beliefs. Finally, section 4 claims that even if rapid, automatic, and unconscious informational integration were characteristic of our common sense conception of beliefs, externally stored 'beliefs' would still be sufficiently similar to beliefs properly so called for them to be grouped together for all practical and scientific purposes.

1 The Extended Mind Thesis and the Causal Role of Beliefs

The extended mind thesis says that "at least some mental states may be partially or wholly constituted by states outside of the brain and body" (Weiskopf 2008, p. 265). When certain features of a cognitive system's environment are appropriately coupled to its internal cognitive processing, they become, literally, a part of it. Cognitive systems are thus not constrained by the physical boundaries of their skin, or skull, but extend into the environment. For instance, since beliefs are characterized by their functional role, and since there is no reason why that role can be played only by something from inside a cognitive system, "beliefs can be constituted partly by features of the environment, when those features play the right sort of role in driving cognitive processes" (Clark and Chalmers 1998, p. 11).

Consider Clark and Chalmers' famous example of Otto, who suffers from Alzheimer's disease and relies on a notebook to help him structure his life. If certain conditions are met—if the notebook is a constant in Otto's life, if its information is directly available to him without difficulty, if, upon retrieving its information, he automatically endorses it, and if its information has been consciously endorsed at some point in the past—then "Otto's internal processes and his notebook constitute a single cognitive system" (Clark and Chalmers 1998, p. 16). In that case, the states of Otto's notebook count as his beliefs, although they are stored not internally, but externally.[ii]

According to Weiskopf, in contrast, the states of Otto's notebook do not count as beliefs properly so called because they aren't usually amendable to rapid, automatic, and unconscious informational integration: "Beliefs are, as I will say, normally informationally integrated with, and updated in concert with, other beliefs …. But most externally located mental states do not share this feature. So … they cannot be beliefs" (Weiskopf 2008, p. 268). Informational integration, he says, "is normally achieved by mechanisms that operate below the threshold of our conscious awareness and control … [and that] typically operate rapidly, allowing many beliefs to be updated in the same relatively brief span of time" (Weiskopf 2008, p. 268). Thus, when a believer acquires a new belief properly so called, her belief system is "automatically and unconsciously updated to reflect this new information" (Weiskopf 2008, p. 269), but not when she acquires an externally stored 'belief,' say a new entry in her notebook:

> Waldo … 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, and sometimes he succeeds in this. Now suppose 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. … 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 this new information about the museum. (Weiskopf 2008, p. 269)

The same holds for 'pure external memory situations' in which an acquired externally stored 'belief' fails to be integrated with other externally stored 'beliefs':

> 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 will be updated or erased to reflect his new total informational state. (Weiskopf 2008, p. 269–270)

Since externally stored 'beliefs' are either not updated at all, or only slowly, non-automatically, and consciously (by, say, checking all one's notebook entries), they do not

count as beliefs properly so called. The extended mind thesis with regard to beliefs, Weiskopf argues, is wrong.[iii]

2 Is Informational Integration Usually Rapid, Automatic, and Unconscious?

Interestingly, Weiskopf has little to say as to why informational integration must usually be rapid, automatic, and unconscious. He acknowledges that "spelling out how informational integration works … relies on a substantially fleshed out empirical theory of human reasoning," which he says he cannot provide. Instead, he offers some examples to illustrate how informational integration "works in quotidian cases" (Weiskopf 2008, p. 268): When an ordinary believer learns that Sam and Max have been divorced, she ceases to believe that they share the same address, and when she learns that the museum has been torn down, she ceases to believe that it is on 53rd Street, that the museum café makes a good latte etc. This, he says, "is a banal illustration of informational integration, but its banality just emphasizes the fact that integration is part of the everyday dynamics of belief" (Weiskopf 2008, p. 268).

Undoubtedly, existing beliefs are revised in the light of new information. What we doubt is that this revision is normally rapid, automatic, and unconscious, or that it must be so in order for informational states to count as beliefs properly so called.

If we acquire a new belief (that the museum has been torn down, say), and if it is remarkable enough for it to be worth memorizing, it is added to our memory. This process of adding a new belief may be automatic and (largely) unconscious. But adding a belief to memory is not yet integrating it with the rest of one's belief system, and

Weiskopf doesn't offer any argument for this further claim (calling it 'banal,' we dare say, doesn't count).

In fact, there is ample reason to be skeptical. We are not isolated 'thinkers' whose minds work independently of their interaction with the world. The inadequacy of such a model, according to which our mind is an information processing device residing in our heads that receives input from the body and the world, does some internal computation, and then produces some output, has been a major topic in recent philosophy of cognitive science (see, e.g., Clark 1997, 2008, Varela et al. 1991). In fact, cognition is intimately tied to agency. We are, as it were, agents, not mere thinkers (for vision, for instance, this point has been forcefully made by Noë 2004 and O'Regan and Noë 2001). In our view, being agents, it is only when we act that the coherence of our belief system matters; unless we act, we can believe all sorts of things. There is thus no reason to suppose that in cases where a new belief is entirely irrelevant for what the believer is doing or planning to do, its acquisition nevertheless automatically and instantaneously triggers a process of informational integration.

Moreover, it seems that we are not constantly updating all the beliefs affected by some newly acquired belief. Upon learning that the museum has been torn down, we do not instantaneously cease to believe that there are exactly seven museums in town, that there is no empty site on 53rd Street between 5th and 6th Avenue, that there is either a museum or a shopping mall on 53rd Street, etc., nor do we come to believe instantaneously that there are now fewer museums in town than last year, that there is now an empty site on 53rd Street, that answering 'Seven' would be a good thing to do when asked how many museums there are in town etc.

Belief revision must always be revision of salient, or relevant, beliefs (this is the lesson of the frame problem). 'Salient,' however, can only mean 'salient with regard to a particular action.' For creatures like us, then, who are not just thinkers, but agents, it will suffice if informational integration is triggered by a concrete situation in which we are about to act. Moreover, for biologically realized cognitive systems with limited resources and capacities, it is the only sensible (computationally or energetically tractable) way informational integration could possibly proceed.[iv] Constantly updating all the beliefs affected by a newly acquired belief would be a Herculean labor (see Clark 2005, p. 7). Nearly every second of our waking life we are forming new perceptual and non-perceptual beliefs, which are related to dozens of others, which in turn are again related to myriad others. If parsimony and efficiency matter, we should not expect the mental economy of cognitive systems to work the way Weiskopf thinks it does—neither for natural cognitive systems that are the result of natural selection, nor for artificial cognitive systems that are the result of a deliberate design process.[v]

Hence, with regard to the rapidity and automaticity of informational integration, Waldo's case is unproblematic. Suppose that having stored belief A, that the museum is on 53rd Street, and belief B, that the museum café sells a good latte, externally, he acquires at t_1 a new belief C, that the museum has been torn down, which he stores internally. The acquisition of C initially leaves A and B unaffected. At t_2, however, Waldo is craving for a good latte and is checking his notebook for latte-related entries. Once he finds the entries 'The museum cafeteria sells a good latte' and 'The museum is on 53rd Street,' he realizes that they conflict with his internally stored belief that the

museum has been torn down and updates them accordingly. Exactly the same, we venture, happens in the case of an ordinary believer who has stored all beliefs internally.[vi]

What about the claim that informational integration is unconscious? Since 'automatic' and 'unconscious' are used interchangeably in some contexts (for instance in the automaticity-research in social psychology done by John Bargh and colleagues; see, e.g., Wegner and Bargh 1998; Bargh and Chartrand 1999), it would have been helpful had Weiskopf clarified what distinction he had in mind. One sensible suggestion is that informational integration is automatic if it is triggered solely by the acquisition of a new belief, while it is unconscious if it is not something that is consciously monitored or controlled by the subject. With 'unconscious' so understood, Waldo indeed seems to be different from an ordinary believer: The latter, when pondering about where to get a good latte, recalls that the museum has been torn down and updates her (situationally salient) beliefs effortlessly, fast, and without conscious monitoring. Waldo, in contrast, has to check all his notebook entries and revise the relevant ones on a piece-by-piece basis, which takes longer and requires conscious effort. Hence, even if informational integration is not usually rapid and automatic, there seems to be an important functional difference between Waldo and ordinary believers with regard to whether or not it normally requires consciousness.

A rather radical line of response would be to deny that informational integration is usually unconscious in ordinary believers. If informational integration is indeed a matter of the believer's being in a particular situation in which she is about to act, and thus happens 'on-the-fly,' one may legitimately question whether it happens without the subject's being consciously aware of it. Suppose you are meeting a friend at the subway

station on 51st Street and Lexington and you decide to go and have a cup of coffee. You believe that the museum is on 53rd Street and that the museum café makes a good latte. At that moment, you recall that you have read in the newspaper last week that the museum has been torn down. As a consequence, you conclude that the museum is no longer the place to go if you want a good latte and revise your beliefs accordingly. It does not seem as if the informational integration happened outside of conscious awareness, and this case is just as ordinary as myriad others.

Since Weiskopf does not say what exactly he means by 'unconscious,' we venture that he may be conflating 'unconscious' with 'effortless' (on pp. 268 and 271 he seems to equate 'conscious' and 'effortful,' which at least suggests that 'unconscious' is supposed to be equivalent to 'effortless'). Ordinary subjects usually may have to put so little effort into informational integration that it may seem as if this process were unconscious, while in fact it is conscious, but lasts only a split of a second, due to its effortlessness. Again, if informational integration is not usually unconscious, then Waldo's case is unproblematic. Of course, there is still a difference between Waldo and ordinary subjects with regard to the amount of effort required for informational integration. That, however, is only a quantitative, and not a qualitative difference, and a quantitative one does not warrant Weiskopf's conclusion of a difference in kind between Waldo's 'beliefs' and the beliefs of an ordinary subject.[vii]

A more moderate line of response is that the proponent of the extended mind thesis is not committed to claiming that Waldo or Otto are indistinguishable from ordinary believers tout court. How could they be, given that their biological long term memory is severely impaired? A hand prosthesis may be slower in operation than or not able to exert

the same amount of pressure as a biological hand, but it is still a hand, albeit an artificial one. Likewise, Clark and Chalmers' initial suggestion concerning Otto was that "taken as a single, integrated system, Otto-and-the-notebook exhibit ENOUGH OF THE CENTRAL FEATURES AND DYNACMIS OF A NORMAL AGENT having (amongst others) the dispositional belief that MOMA is on 53rd St. to warrant treating him as such" (Clark 2005, p. 7; long capitalization M.K. & S.W.). Otto thus is a 'true believer,' albeit different from ordinary believers.

It could be objected that while the exact grasping behavior is not part of the functional profile of a hand, it is part of the functional profile of belief that new beliefs are usually integrated unconsciously (or effortlessly). To that issue we will return in section 3. Before that, we will briefly address two other possible objections.

According to Weiskopf, we commit ourselves to an untenable inconsistency and irrationality of believers with externally stored beliefs. The problem, as he puts it, is that a believer acquiring a new externally stored belief B which is in conflict with a belief A that's also stored externally will typically continue to believe A:

> [I]f we suppose these external states [the notebook entries; M.K. & S.W.] 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. (Weiskopf 2008, p. 270)

Suppose Otto's notebook reads 'The museum is on 53rd Street' on page 10 and 'The museum is no longer on 53rd Street' on page 26. Does that render Otto irrational? We don't think so. One strategy would be to accept that there is a sense in which Otto has both beliefs, but to insist that merely having both beliefs does not render him irrational as long as he does not consciously endorse them simultaneously. Irrationality, one could put it, is considered endorsement of conflicting beliefs in the light of prospective action. The fact that the beliefs are in conflict with each other is irrelevant as long as Otto's actions are not guided by both of them. He would be irrational, or rather act irrationally, if, upon reading on page 26 that the museum is no longer on 53rd Street, he nevertheless kept answering 'Yes' when asked whether the museum is on 53rd Street and went to 53rd Street to get a latte. But that is not what is going to happen in typical Otto-and-the-notebook cases in which Otto carefully checks his notebook entries before acting.[viii]

An alternative response would be that no matter what is written in Otto's notebook, he has only one of the two beliefs in question, depending on which one is guiding his actions (such a response would square nicely with an interpretationist take on beliefs and belief ascriptions along the lines of Dennett 1971, 1987).

Another objection may be that we are conflating standing and occurrent beliefs. Weiskopf is explicitly talking about a subject's standing beliefs, which are "the contents of some portion of her long term memory store" (Weiskopf 2008, p. 269). While standing beliefs are usually updated rapidly, automatically, and unconsciously, it could be said, the beliefs we have been talking about that are not updated in such a way and are guiding our actions are occurrent beliefs. Yet, there is no reason to think that Otto's or Waldo's

beliefs that the museum is on 53rd Street, before or after the informational integration, are not standing beliefs and that the informational integration prior to action does not concern their standing beliefs. On our suggestion, the subject's attention is situationally directed toward an action-relevant belief, and this may create the impression that we are talking about occurrent beliefs. In fact, however, this action-relevant belief is a standing belief, viz., the contents of the subject's long term memory store—the contents of the notebook in Otto's and Waldo's case, the contents of our biological memory system in our case.

Our alternative account of informational integration is at least as plausible as Weiskopf's, and it is compatible with externally stored beliefs' being beliefs properly so called. We now consider what would happen if we are wrong and informational integration in ordinary subjects is indeed typically rapid, automatic, and unconscious.

3 Informational Integration and the Functional Role of Belief

Suppose that brain-stored beliefs are usually updated rapidly, automatically, and unconsciously. Does it follow that external states that are not integrated in such a way cannot be genuine beliefs? Imagine we encounter an alien species the members of which, for all we can tell, are as intelligent and rational as we are. However, also suppose that eventually we find that their memory/belief system (be it biological or purely 'digital') usually relies on delayed, non-automatic, conscious, and effortful informational integration.[ix] Should we seriously conclude, after years of successful interaction, that the members of this species, all appearances to the contrary, had never really believed anything, and that, if we dare say that they ever really believed something, we are thereby

committed to calling them inconsistent and irrational? The fact that the answer to this rhetorical question can only be 'No!' (for some interesting and compelling reasons why see, e.g., Dennett 1971, 1987) shows that the rapid, automatic, and unconscious nature of informational integration is not part of the functional profile of belief (see Clark 2005, p. 6).

Weiskopf anticipates the objection "that informational integration is not a necessary feature of belief" (Weiskopf 2008, p. 273). In response, he argues that rapid, automatic, and unconscious information integration (in typical cases) is a necessary feature of beliefs properly so called because it is required for predictive purposes: If someone tells you that the museum has been torn down, you expect her to not to be going there any more; but since externally stored 'beliefs' are not updated, he ventures, no such useful predictions could be made for believers with externally stored beliefs (Weiskopf 2008, p. 273–274).

To see why this is beside the point, note that predictions of beliefs are only relevant because and insofar as they inform us about future actions. Suppose we tell you that we have moved to a different state. What Weiskopf seems to be saying is that based on this we can predict a change of your beliefs—you no longer believe that we have the same phone number—which is not what would happen were you storing your beliefs externally, since the state of your address book would not change automatically. Yet, what we predict you believe matter only insofar as it affects what we predict you will do. Predicting your beliefs per se, in contrast to predicting your actions in a given situation, is pointless, since such predictions are neither independently testable, nor explanatory. Our crucial concern is: Will you still call us at our old place when inviting us to your wedding, or will you call us at our new place? In typical cases, we predict that since you

want to call us, you will look up our number in your address book, and that this will make you recall that we've told you that we've moved to another state (or you notice, when dialing the old number, that it is no longer ours). As a consequence, you will call the directory assistance to enquire about our current telephone number, as a result of which you will (hopefully) be able to give us a call at our new place.

Weiskopf is right that informational integration is indispensable from the point of view of our beliefs' predictive function. But that is a far cry from showing that it must usually be rapid, automatic, and unconscious. We have argued above that informational integration is a situation-dependent, action-oriented 'on-the-fly' update, and that this is entirely compatible with the view that beliefs serve an important explanatory and predictive purpose, as long as we keep in mind that predicting beliefs is only a means to the end of predicting situated behavior.

The functional profile of belief hence does not require that newly acquired beliefs must usually be integrated with already existing beliefs rapidly, automatically, and unconsciously.

4 Does it Matter, After All?

The foregoing considerations should suffice to drive home our point that the extended mind thesis as applied to beliefs has nothing to fear from Weiskopf's considerations concerning the functional role of belief (for further discussion of this issue see, e.g., Chalmers 2008; Clark 2005; Gertler 2007). Suppose, however, that rapid, automatic, and unconscious informational integration is both de facto characteristic of ordinary human

believers and an essential aspect of the functional profile of beliefs properly so called. Does it follow that an ordinary believer and the Otto-and-the-notebook system have nothing interesting in common, nothing, that is, that would justify treating them as similar kinds of cognitive systems for the purpose of everyday life and psychological science? We don't think so. Even if the ultimate verdict is that an ordinary believer has beliefs properly so called, while Otto has only schmeliefs, the former and the latter are significantly similar for it to be warranted to treat them as one and the same phenomenon. As Chalmers (2008, p. 000) puts it:

> Even if commonsense psychology marks a distinction here, the question still arises of whether this is an important distinction that ought to be marked in this way. One can argue … that Otto's extended state involving the notebook functions in explanation in very much the way that beliefs function in psychological explanation. If so, then it ought to be classified as a belief, whether or not it is so classified by common sense.

As we have argued in sections 2 and 3, there seems to be no significant difference between Waldo, Otto, and ordinary human beings as far as their rationality, their behavior, the explainability, predictability and intelligibility of their behavior and the cognitive processes that precede their behavior are concerned.

Compare the following two dialogues: A: "Do you happen to know Andy's phone number?" – B: "Yes, I do. Hold on a second! I just have to boot my blackberry." versus: A: "Do you happen to know Andy's phone number?" – C: "Yes, I do. Hold on a second!

I can't think of it right now, ask me again in a couple of minutes." There is a perfectly straightforward sense of 'know' in which both B and C are right in their claim that they know Andy's phone number. There is an important commonality between B and C: both have a constant source of information (the blackberry in B's case, we may suppose, the brain in C's case) which usually is directly available to them, and which is such that upon retrieving information from this source, they automatically endorse it, where the information is there because it has been consciously endorsed at some point in the past. For all practical and psychological purposes, that renders B and C alike. If you need to know Andy's number, asking C is as good as (if not better than) asking B; predicting that C will be asking for an international call when trying to call Andy from New York is as reliable as predicting that B will do so; considering it a lie when B tells the police she does not know Andy's number is as justified as considering it a lie in C's case; and so on. If you keep insisting that B really knows it, while C only sort of knows it, then so be it. In our eyes, however, that's a difference without a difference.

5 Conclusion

According to Weiskopf, externally stored extended 'beliefs' fail to be beliefs properly so called because they are not usually integrated with a believer's already existing beliefs rapidly, automatically, and unconsciously. We have argued that none of these features is exhibited by the brain-stored beliefs of ordinary believers, and that they fail to be definitive of beliefs, or their explanatory and predictive function. We have suggested an alternative account of informational integration that stresses the situated and action-orientated character of informational integration which is compatible with extended

beliefs and allows us to preserve central intuitions concerning rationality and belief-based predictions of actions.

References

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

Adams, F. & Aizawa, K. (2008). The bounds of cognition. (Oxford : Wiley-Blackwell)

Ballard, D., Hayhoe, M., Pook, P., & Rao, R. (1997). Deictic codes for the embodiment of cognition. Behavioral and Brain Sciences, 20, 723–767.

Bargh, J.A., & Chartrand, T.L. (1999). The unbearable automaticity of being. American Psychologist, 54, 462–479.

Chalmers, D. (2008). Foreword to Andy Clark's Supersizing the Mind. (In A. Clark, Supersizing the mind: Embodiment, action, and cognitive extension (pp. 000-000). Oxford: Oxford University Press.)

Clark, A. (1997). Being there: Putting brain, body and world together again. (Cambridge, MA: MIT Press)

Clark, A. (1998). Magic words: How language augments human computation. (In P. Carruthers & J. Boucher (Eds.), Language and thought: Interdisciplinary themes (pp. 162–183). Cambridge : Cambridge University Press.)

Clark, A. (2005). Intrinsic content, active memory and the extended mind. Analysis, 65, 1–11.

Clark, A. (2008). Supersizing the mind: Embodiment, action, and cognitive extension. (Oxford: Oxford University Press)

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

Dennett, D.C. (1971). Intentional systems. Journal of Philosophy, 68, 87–106.

Dennett, D.C. (1987). True believers. (In: The intentional stance, (pp. 13–35). Cambridge, MA: MIT Press.)

Gallagher, S. (2005). How the body shapes the mind. (Oxford: Oxford University Press)

Gertler, B. (2007). Overextending the mind. (In B. Gertler & L. Shapiro (Eds.), Arguing about the mind (pp. 192–213). London: Routledge.)

Gibbs, R. (2005). Embodiment and cognitive science. (Oxford: Oxford University Press)

Huxley, T.H. (1895). Collected essays (vol. 3): Science and education. London: Macmillan and Co.

Noë, A. (2004). Action in perception. (Cambridge, MA: MIT Press)

O'Regan, K. and Noë, A. (2001). A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences, 24, 939–1031

Rupert, R. (2004). Challenges to the hypothesis of extended cognition. Journal of Philosophy, 101, 389–428.

Varela, F., Thompson, E. & Rosch, E. (1991). The embodied mind. (Cambridge, MA: MIT Press)

Wegner, D.M. & Bargh, J.A. (1998). Control and automaticity in social life. (In D. Gilbert, S. Fiske, & G. Lindzey (Eds.), Handbook of social psychology (pp. 446–496). Boston: McGraw-Hill.)

Weiskopf, D. (2008). Patrolling the mind's boundaries. Erkenntnis, 68, 265–276.

[i] If Clark and Chalmers are right, it may be impossible to fully understand cognitive systems without taking into account their bodies (acknowledging the essentially

embodied nature of cognition), their environment (acknowledging the essentially situated nature of cognition), and their dynamical interaction with the environment (acknowledging the essentially enacted nature of cognition)—ideas that have played an important role in recent philosophy and cognitive science (see, e.g., Gallagher 2005, Gibbs 2005, Varela et al. 1991).

[ii] Clark and Chalmers do not maintain that these conditions are individually necessary or jointly sufficient, and they seem willing to consider abandoning the fourth condition (conscious endorsement in the past) because it would seem to include beliefs acquired via subliminal perception or memory tampering (Clark and Chalmers 1998, p. 18).

[iii] Weiskopf (2008, p. 271) acknowledges that there may be extended beliefs properly so called, if, for instance, they are stored in some kind of portable electronic memory device that is directly jacked into a subject's brain and does the quick and automatic updating required by our common sense conception of belief. His point is that given present technology, there are no extended minds.

[iv] Weiskopf (2008, p. 268) says he cannot provide an "empirical theory of human reasoning" that is "saying under what circumstances a set of states is relevantly related to another." That's a pity. Realizing that the relevant beliefs are those that are relevant for action allows one to see why informational integration need not be rapid and automatic.

[v] Consider a study by Ballard et al. (1997), discussed in Clark (2008, p. 000). The subjects' task was to copy a pattern of colored blocks on a computer screen by using the mouse to move colored blocks from a reserve area to a new workspace, one at a time. From a 'pure-thinker' perspective, the strategy apparently would be to look at the target, memorize the color and position of the block to be added, and then get an appropriate

block from the reserve area (which is indeed what subjects report). Using eye tracking technology, however, Ballard et al. discovered repeated rapid saccades to the model both before and after picking up a block, suggesting that the subject memorizes only one piece of information (color or location) at a time, while repeated fixations provide the information 'on the fly' when it is required by the action to be taken (picking up a block of a certain color vs. placing it at a particular spot).

[vi] Note that by distinguishing Waldo from 'ordinary believers' we do not mean to suggest that Waldo or Otto are 'unnatural' or 'abnormal.' On our view, extended cognitive systems like Waldo or Otto, although they may still be exceptional, are decidedly not unnatural.

[vii] Note that sometimes informational integration requires conscious effort even in ordinary subjects, e.g., when we are thinking through a philosophical argument, trying to find out which assumption to give up or modify in order to avoid contradiction. Since the conscious/unconscious distinction arguably doesn't admit of degrees, Weiskopf would have to hold that such cases are clearly separable and distinct in kind from ordinary unconscious cases, which seems counterintuitive. On our view, in contrast, the difference would not be a qualitative but only a quantitative one along a continuum of cases which differ with regard to the amount of cognitive effort required by informational integration.

[viii] Storing part of one's beliefs externally may even increase the chance of one's making the right decision, since purely internal belief systems are not always reliable: we forget things, misremember them, and are prone to failures of rationality when relying solely on our internal cognitive system (as revealed by dozens of studies on human decision making in social psychology). As a matter of fact, the literature on the embodied nature

of our minds is full with examples to the extent that 'outsourcing' cognitive tasks enhances a subject's cognitive capacities (see, e.g., Clark 1997, 1998). Although extended cognitive systems like Waldo and Otto might still be exceptional, systems that extend their belief systems by partially relying on external tools seem to become more common rapidly.

[ix] We may assume that they are able to process more information consciously per time unit than we are so that we don't realize, just by talking to them, that their way of informational integration differs radically from ours.