

Gregory Nirshberg
University of Texas at El Paso
gnirshberg@miners.utep.edu

It's Not All About Beliefs: The Extended Mind and Why Propositional Attitudes are Still in the Head

Introduction

The extended mind hypothesis (henceforth EMH) says that cognition, and thus our minds, extends beyond the boundaries of our skin and skulls. Proponents of EMH mean this to be a claim about cognition in general, a claim that should hold for all of our propositional attitudes, but the claims thus far have focused almost exclusively on beliefs. I claim that the neglect of other propositional attitudes poses a significant problem for proponents of EMH. Through a series of thought experiments, I'll show that even if EMH ends up being true, it is at best a claim about extended beliefs, not extended cognition. These examples will further show that much of the appeal of EMH rests on a conflation of propositional attitudes with propositional content, a confusion only made possible by the focus on belief.

The Extended Mind Hypothesis

Initially formulated by Andy Clark and David Chalmers¹, the most popular thought experiment used to support EMH is the case of Otto, a man with Alzheimer's, who carries around a notebook in which he writes down all his beliefs, and to which he refers regularly to aid him in his cognitive endeavors (12). Otto is contrasted with Inga, a woman with normal cognitive functioning. Both Otto and Inga want to go to the Museum of Modern Art. Inga thinks for a moment, recalls it's on 53rd street, and then goes there. It's clear that Inga believed that the museum was on 53rd street. Otto, on the other

¹ Though the original formulation of the extended mind hypothesis was made jointly by Clark and Chalmers, of the two, Clark has gone on to be the central proponent and defender of the view. And though others have also defended similar claims, Mark Rowlands for example, in this paper I will be directing my criticisms explicitly towards Clark. Rowlands, in particular, explicitly claims that he takes EMH to be primarily about mental processes, not mental states (629), and finds the case of Otto problematic as an exemplar (631). This paper will have a particular focus on an unexamined issue having to do with propositional attitudes, though it should be noted that much of this critique should be able to generalize to an account such as Rowlands, in so far as our propositional attitudes play a constitutive role in our mental processes.

hand, checks his notebook, sees that the museum is on 53rd street, and then also goes there. Clark argues that since Otto's notebook plays the same functional role that it would, had all the information in it been in Otto's head, then the information in it should be counted as a belief, available to Otto's cognitive processes.² Otto believed, before he checked the notebook, that the museum was on 53rd street³.

Let us grant Clark his desired conclusion regarding the case of Otto: Otto has a belief about the Museum of Modern Art's location, stored in the notebook. But EMH isn't exclusively framed as a claim about extended beliefs; it is put forward as a claim about the *mind* extending beyond the skin and skill boundary (Clark and Chalmers 12). EMH is then, at least in principle, a claim about all mental states (Clark 964). Clark's focus on beliefs is not surprising: it is standard for philosophers of mind to identify beliefs and desires as the paradigmatic propositional attitudes and, even then, to focus almost exclusively on beliefs.⁴ Despite this focus, there is tacit agreement that claims about our propositional attitudes should include the full range of attitudes – not only beliefs and desires, but also hopes, fears, etc.^{5,6} The exclusive focus on beliefs has left proponents of EMH the task of making good on the broader claim in regards to the rest of the propositional attitudes.

Clark makes a passing reference to the possibility of extended desires: a waiter at your favorite restaurant who has information about your favorite meal might be said to store your desire for that meal (17-18), but the example is not developed any further. In his critique of extended cognition, Dan

² Clark codifies this in what he labels the parity principle: "If, as we confront some task, a part of the world functions as a process which, were it to go on in the head, we would have no hesitation in accepting as part of the cognitive process, then the part of the world is (for that time) part of the cognitive process." (Clark 2)

³ Clark and Chalmers do specify that certain constraints and satisfaction conditions must be placed on the information to count as a belief, constraints such as that the information is reliably accessible, available to consciousness, and available to guide action (13). The information also had to have been consciously endorsed at some time in the past and is automatically endorsed upon being retrieved (Clark and Chalmers 17).

⁴ See Stich 573-591

⁵ See Churchland 71, for more explicit enumeration.

⁶ Even something as important as intentions have been absent from extended mind literature.

Weiskopf considers how EMH would handle the integration of extended beliefs with other mental states. He argues that an *external* belief cannot integrate with an *internal* desire in the proper way to affect the behavior of the agent (275). But as of yet, no one has taken the move to show how desire itself can be extended. I'll show that EMH, in principle, cannot deal with these other propositional attitudes.

Extended Desires

Let us imagine Martha. Martha is in a similar situation to Otto (in terms of having Alzheimer's), except Martha also keeps some of her desires housed in her notebook. Martha has become unhappy with her weight and has decided she would like to lose 15 pounds. She proceeds to write down her desire to lose 15 pounds in her notebook. She further decides that to reach the satisfaction conditions of her desire she must do the following: work out at the gym, change her eating habits (which involves cutting sweets and fried foods out of her diet), walk to work, and more. Imagine it is a couple days later and Martha wakes up one morning. She doesn't pack her gym clothes. She walks out of her house and gets in her car and drives to work. During her lunch break a coworker offers her a brownie and she takes it without hesitation. Later that night she watches a movie while munching on potato chips. Before bed she looks at herself in the mirror and admires her figure, "I'm one sexy lady," she thinks to herself.

Given that it was stored in the notebook, not only was the desire unable to play the proper role in influencing Martha's behavior, but it seems queer to even ascribe to Martha the desire to lose 15 pounds at all. We can imagine the struggle of Samantha, a woman without Alzheimer's, who also decides to lose 15 pounds. She wakes up in the morning and depressingly gets together her gym clothes. She leaves the house and longingly looks at her car as she walks down the street, frustrated, but envisioning a skinnier version of herself months in the future. During lunch, her coworker offers her a cupcake, and the internal battle she wages to turn it down is rivaled only by drug users and gamblers in

the later stages of a debilitating addiction. Later that night, before bed, she spends a minute looking at herself in the mirror. Tired and exhausted from a long day (and a tough workout at the gym), she notices no discernible loss of weight. She knows changes don't happen overnight, but boy, sometimes she wishes that she didn't have this stupid desire to lose weight, and could go to sleep without these conflicting emotions. She falls asleep craving cupcakes, but, also, proud of her resolve earlier that day.

It might be argued that this example misses the mark, since the purpose of the notebook that Martha carries is that it is available to her when she needs it. This is in fact the very reason that Clark stressed the criterion that the information be reliably available to consciousness when needed to guide action (13). The information *was* available; Martha was just not taking advantage of it. But this is precisely the point. Long standing desires are not propositions that we access from time to time when they are needed to guide action (such as the address of the museum in the Otto example); attitudes such as these play a consistent and regular role in our behavior and mental lives. It is unrealistic for Martha to check her notebook constantly throughout the day to see which, if any, of her extended desires might be relevant in directing her behavior at the present moment; but this seems to be the only way to ensure that Martha acts in accordance with her desire to lose 15 pounds. The role that we expect a desire to play in our behavior doesn't seem to be one that can function properly when offloaded to an external store; unlike beliefs, we don't stop and introspect to see if we do or do not desire something before acting.⁷

A natural response might be to say, "We all need reminders of our goals from time to time; things to keep us on track. Why is the content in Martha's notebook any different from a post-it note placed on the fridge reminding us not to indulge in treats?" While intuitively appealing, this rejoinder

⁷ It is true that we stop and examine ourselves for the presence of a short term desire from time to time, "do I want strawberry or chocolate ice cream right now?" But our long term standing desires tend to be significantly more functionally integrated in our behavior. The example in the next section on extended fears will develop this point further.

fails as well. Imagine that Samantha from above places a post-it note on her fridge which says “DON’T EAT THE CUPCAKE!” Maybe Samantha was engrossed in a good book, got hungry, and was about to grab a snack from the fridge when she saw the note. She stops, struggles for a moment, damns her diet, and grabs a banana from the counter instead. While it’s true the note played the proper functional role in Samantha’s behavior, it’s more natural to look at the note as a cue which engaged Samantha’s internal desire to lose weight. In fact, the use of the note *only* makes sense if there is an internal desire already present, allowing Samantha to interact with the note in the proper way.

Imagine the same situation with Martha. When Martha reads the note telling her not to eat any cupcakes, it is unclear why she should act in light of that information. Why not just disregard it? She’s not aware of any desire to not eat cupcakes. This is because the note itself is not the desire, but must integrate with an already existing desire in the proper way. At best the note could be said to externalize the content of the desire, but not the attitude itself; a point which I will return to later. Alternatively, it might be argued that in Martha’s case the note should read “don’t eat the cupcake because you desire to lose 15 pounds.” But this explicit formulation of the desire should be just as worrisome to Clark. It is not the case that Otto writes in his notebook, “I believe the museum is on 53rd street.” The fact that Otto believes, rather than hopes or fears, the proposition is apparently implicit in the storage of the information, something that doesn’t seem to transfer in the proper way to other propositional attitudes. This is why the case of Martha is so problematic. The fact that Martha would struggle with her craving to eat a brownie, or be pleased or upset with her appearance in the mirror, based solely on whether or not she happens to be reading her notebook at the time, strikes us as odd⁸. The plausibility of extended attitudes becomes even more pressed when we venture further beyond the standard propositional attitudes of belief and desire to consider states like fear, which I turn to in the next section.

⁸ Unless we want to posit that Martha has a standing desire to desire whatever the notebook indicates she desires. But Clark seems to want to specifically avoid this sort of conclusion (Clark and Chalmers 13).

Extended Fears

Imagine Frank. Frank is similar to Otto, except Frank also keeps some other propositional attitudes stored in his notebook. One day Frank is sitting in his study working at his desk when he notices a large spider crawling on the wall next to him. “What an interesting looking creature,” Frank thinks to himself, “Spiders have such an odd way of moving.” He gets up to take a closer look. As is his habit, Frank also checks his notebook to see if he has any relevant beliefs or desires in regards to spiders. When he gets to the proper section he reads the following, “You have severe arachnophobia. Spiders terrify you.” Immediately Frank falls back. His heart is racing. His hair is standing on end. He finds himself choked up, unable to speak. He is also frozen in place, unable to move except for the shaking of his hands. Finally Frank musters enough strength to reach for his notebook, and continues reading on the page, “Don’t worry though, you went through immersion therapy and no longer fear spiders.” Suddenly he starts to feel better. His heart rate slows down, he stops shaking, he can get up and look at the spider again. “Phew,” he thinks, “I should really delete that part of the book so this doesn’t happen again!”

The above scenario should strike the reader as highly implausible and counter-intuitive. We can run the same thought experiment with other fears, or with hopes and wishes, but EMH doesn’t seem to work for these propositional attitudes. The kind of updating and integration we expect from them simply doesn’t make sense with the content in the notebook. It is not only that our folk intuitions would be troubled by someone who claimed to keep all their fears and hopes and desires in a notebook, and that we would be unlikely to consider the agent the genuine owner of those propositional attitudes if they did, but that functionally, there seems to be no way to flesh out how this would work.

Beliefs and Cognitive Processes

Clark might try to respond by shifting the focus to beliefs, and the way those beliefs integrate more broadly with our cognitive processes. So, for example, when I am sitting at my desk doing math problems, or working through a logic puzzle, or doing some bit of analytical reasoning, the entire process of interacting with the external belief is how we flesh out the import of EMH. But even granting the importance of beliefs, these other propositional attitudes are equally, if not more, important.

Individuating a cognitive system requires more than an understanding of the beliefs that an agent holds, and more than an understanding of the role those beliefs play in various cognitive functioning. It also requires understanding how all those beliefs come together and integrate with their hopes and fears and desires and intentions. We can posit two individuals with an incredible amount of overlap in their beliefs, and even cognitive skills, but what we take to be relevant about them is how those agents behave given the way in which those beliefs interact with their other mental processes. Clark explicitly points out the importance of extended beliefs being ‘appropriately poised’ to guide behavior (967), but has thus far neglected the (likely more primary) role of our other propositional attitudes in determining our behavior. Whether I give up in frustration after working on a math problem for five minutes, and whether you continue to pluck away at it for an hour, says something very important about our respective minds, with very little regard for what beliefs we hold⁹. Imagine for a moment that you and I both come to hold the belief that god doesn’t exist. Whereas I hoped for this proposition to be true, you feared the truth of this proposition. I suddenly have a new spring in my step and newfound sense of freedom. You are undergoing a deep dark existential crisis and have been unable to drag yourself out of bed for the last week. What is interesting about the minds of these two agents is

⁹ One attempted solution would be to posit that one of us has certain beliefs regarding the value of not giving up, and so on, but standard arguments surrounding motivation will demand another desire be present to explain why the agent behaves in light of that belief; since we can posit that both agents have that same belief, but only one desires, in the proper way, to act in light of it.

not the shared belief that god doesn't exist, but how that belief integrates with the rest of their mental states, which, if my above examples are sound, remain firmly grounded within the classic skin/skull boundary. Earlier I granted Clark the claim about beliefs, but EMH becomes quite constrained if that is all we're able to grant.

Propositional Attitudes vs. Propositional Content

My final point is a bit more diagnostic about EMH. The previous considerations do more than constrain Clark's claims; they also shed some light on an important, yet neglected, distinction between our propositional attitudes and the content of those attitudes. This distinction tends to go unnoticed because of the seemingly transparent relationship between propositional content, and taking the attitude of belief towards that propositional content. The act of storing content and believing that content seem to be very similar, if not the same thing, because to believe something is just to endorse that content. But storing content and desiring it or fearing it doesn't lend itself to the same type of analysis.¹⁰ Because of this feature of propositional attitudes that is particular to beliefs, it is only the exclusive focus on them that has allowed EMH to work, in so far as it has, with this conflation going unnoticed.

Clark wants to talk about certain standing dispositions that Otto has to the content of his notebook (Clark 5, 7), and that Inga has towards the content in her memory, but it's bizarre to claim this about anything that is not a belief. All our other propositional attitudes seem to have a more nuanced relationship with their content. To modify the EMH claim to account for this distinction between propositional attitudes and propositional content, Clark would have to say that only the propositional

¹⁰ This conflation is likely to be relevant well outside of debates in the extended mind literature. And if the reader at this point is convinced that we should get rid of the entire propositional attitude framework, or the paradigm of content storage in philosophy of mind, I'm sympathetic. The above considerations might be all the more reason to abandon these approaches. But this isn't further motivation for Clark's case at all, since his view depends in large part on this traditional structure.

content of the belief is extended, and that the attitude taken towards it remains firmly in the head. This would be an unacceptable modification, ceding far too much ground, since beliefs, as we commonly understand them, would no longer be extended. A claim that sometimes the propositional content of our attitudes is extended into the environment is far less interesting and radical a claim than originally put forth.¹¹

Conclusion

This is just one of many criticisms that have been made that also serve to overly restrict the import of Clark's claim.¹² What the above cases I considered make clear is that even if the content of an attitude is externalized, the attitude itself never is. If an agent had to explicitly write "I desire..." before the content of all her desires, it's unlikely we'd want to say she actually had those desires, let alone why she herself would want to endorse such a desire when reading it in a book. But if we reflect for a moment on this, it becomes ambiguous as to what Clark ever intended for other propositional attitudes. He assumes that "the MOMA is on 53rd St." written in Otto's notebook is a belief, by, I assume, adding in an implicit "I believe" before it. But it's also possible that Otto fears the MOMA is on 53rd St., or hopes the MOMA is on 53rd St. etc. This would entail that the content of the notebook can only play the right

¹¹ Clark notes that his claim is one about the material vehicles of mental states and processes, not about their contents (Clark 966). It's important to note that the discussion in the current essay about the distinction between the content of propositions and our propositional attitudes is orthogonal to Clark's distinction. It would still remain the case that, at best, the material vehicles of cognition could only account for extended propositional content, not extended propositional attitudes.

¹² This includes criticisms by Robert Rupert, Daniel Weiskopf, and Adams & Aizawa, among others. Responding to the claim that seeing cognition as extended would make a significant difference in scientific investigation (Clark and Chalmers 10), Rupert has argued that Clark's claim doesn't have any significant explanatory import in the cognitive sciences that can't already be gained by positing the much less radical embedded mind thesis (408, 424). Weiskopf also criticizes the claim for an inability to account for the informational integration of extended beliefs, i.e. – that the extended beliefs (at least as described in Otto's notebook) don't play the right sort of functional role (273). It is worth noting that Weiskopf doesn't argue against EMH as a modal claim, but strictly against the claim as one that describes some sort of commonplace phenomenon. Adams and Aizawa argue that, as a matter of contingent empirical fact, cognition is brain bound (46). In part, their argument depends on the fact that the symbols in Otto's notebook have derived content, whereas the mark of the cognitive is non-derived content (Adams and Aizawa 55).

function role when connecting up to an internal attitude of Otto, belief or otherwise¹³. If this the claim Clark is defending, it is no longer the radical claim it has been taken to be. At no point do we really lose the person at the center of all this: a mind that is still constrained by the boundaries of the skull and/or skin.

¹³ Clark wants to avoid saying that what Otto has is a standing belief that the museum is at the location written in his notebook (13), but it seems that this is the most generous thing we can say about Otto in this situation.

Bibliography

- Churchland, Paul M. "Eliminative materialism and the propositional attitudes." *Journal of Philosophy* 78. February (1981):67-90. Print.
- Clark, Andy., Chalmers, David J. "The extended mind." *Analysis* 58.1 (1998):7-19. Print.
- Clark, Andy. "Intrinsic content, active memory, and the extended mind." *Analysis* 65.285 (2005):1-11. Print.
- Clark, Andy. "Spreading the joy? Why the machinery of consciousness is (probably) still in the head." *Mind* 118.472 (2009):963-993. Print.
- Rowlands, Mark. "The extended mind." *Zygon* 44.3 (2009):628-641. Print.
- Rupert, Robert D. "Challenges to the hypothesis of extended cognition." *Journal of Philosophy* 101.8 (2004):389-428. Print.
- Stich, Stephen P. "Autonomous psychology and the belief/desire thesis." *The Monist* 61. October (1978):573-91. Print.
- Weiskopf, Daniel A. "Patrolling the mind's boundaries." *Erkenntnis* 68.2 (2008):265-276. Print.