



Author

ROBERT D. RUPERT

Cognitive Systems and the Extended Mind

Robert D. Rupert, *Cognitive Systems and the Extended Mind*, Oxford UP, 2009, 268pp., \$55.00 (hbk), ISBN 9780195379457. □

Reviewed by Robert A. Wilson, University of Alberta

Once upon a time, empirically-informed, philosophical work on the mind was pretty straightforward. Mental activity went on inside the head, and we were pretty sure that it, along with all the good stuff associated with it -- consciousness, intentionality, mental representation, computation -- could be most readily understood without having to crease our brows too much about how minds were situated vis-à-vis bodies, or vis-à-vis environments.

To be sure, the cognitive sciences had renegade, pluralistic strands running through them, and particular disciplinary itches that needed scratching. But there was a sufficiently articulated paradigm in place that at least made designing jacket covers for books in the field relatively easy. Find a picture of a head, such as a representation of a brain, ideally an image a bit attuned to the title or theme of the book, and you're at least well on your way. I failed to understand this simple point during what were politely called "ongoing discussions" over the cover of *The MIT Encyclopedia of the Cognitive Sciences* a dozen years ago. "No, I'm sorry, Professor Wilson", I was informed by way of a solemn conclusion delivered by an exasperated senior editor at the Press, someone whom I imagined had to avert his gaze in embarrassment on my behalf as he typed back curtly, "I am afraid that it has to be a head".

Designing jacket covers for books in this neck of the woods is no longer so easy. Heads are not exactly out, but they are no longer strictly required. Shaun Gallagher and Dan Zahavi's *The Phenomenological Mind* (2008) has a person half-way through exit mode on its cover. Fred Adams and Ken Aizawa's *The Bounds of Cognition* (2008), for crying out loud, has *a cube of green cubes on the cover*. Not even *a thing that thinks*, so far as we know. Rob Rupert's stimulating book continues down this path with quite a beautiful black and white photograph of a grove of aspens in the wilds of Colorado plonked right on the cover. It's enough to start me worrying that in order to understand the field that I work in, I can no longer simply look at jacket covers.

In order to write this review in my current state of existential confusion, I have thus had to resort to an old trick of the trade, something I learned in grad school but like most things so learnt, haven't had to use much since. Read. And, somewhat to my initial surprise, especially given the total absence of pictures once one gets beyond the cover of *Cognitive Systems and the Extended Mind*, and the distinct dearth of bad jokes between

the sheets, I'm glad I did. Rupert's book is a good read. It is a sustained, systematic, critical examination of the idea that minds are not simply ensconced inside heads, but extend into both bodies and the world beyond the body. As his title suggests, Rupert is focused primarily on the latter of these, the extended mind thesis, a thesis articulated neatly by Andy Clark and David Chalmers a dozen years ago in a paper that, had it been a movie, would have been an instant blockbuster, then a classic, and now be competing with Seinfeld re-runs on cable tv.

Cognitive Systems and the Extended Mind draws on a series of papers published over the past ten years that have been influential, but it is much more than the sum of these past parts. Bookended by introductory and concluding chapters, its ten chapters fall into three constituent parts. In Part I, Rupert defends the idea that "the relatively durable cognitive system ... provides the most plausible line of demarcation between what is cognitive and what is not" (p.7), using this idea to make a prima facie case for the claim that cognition is at least organism-bound. Part II concentrates on arguments for the hypothesis of extended cognition (HEC), and why they fail, with Rupert helping himself to his conclusions from Part I. Finally, Part III moves back from the rejection of HEC to focus on two other ways to understand the sense in which cognition requires some kind of situated approach, what Rupert calls the *embedded view*, and the idea that cognition is *embodied*. I will concentrate on Part I in what follows, say something more briefly about Part II, and leave Part III alone altogether. The bottom line for Rupert: there is no reason to accept the idea of extended cognition, many perfectly sane reasons to reject it, and the take-home lessons from considering the seemingly less radical ideas that cognition is embedded and embodied are that what insights they contain are readily accommodated by cognitive science, traditionally conceived. Paradigm found.

There is much to admire in this book. It is well-structured and well-written, adopting a self-consciously naturalistic perspective on how to understand the mind -- through our best, even if imperfect, empirical sciences in the domain of cognition. By presenting and critiquing a number of explicit arguments for and against the specific views that Rupert considers, *Cognitive Systems* advances the field. For all that, I think that the claims at the heart of the book -- those about cognitive systems that drive the positive argument of Part I, and those that appeal to natural kinds (ch.5), developmental systems (ch.6), dynamical systems (ch.7), and experience (ch.8) in an attempt to undermine arguments for HEC in the negative argument of Part II -- are mistaken. Or, rather, insofar as they are true, they do little to capitalize on the head start advantage that head-bound views of cognition have in virtue of the individualism pervasive in cognitive science's history.

Consider Rupert's chief pair of arguments in Part I. The first of these appeals to an idea central to a number of critiques of HEC, especially those developed by Adams and Aizawa over the past ten years: that there must be a principled way to distinguish "genuinely cognitive contributions" from "merely causal" contributions to cognitive processes (p.16). The idea, in one version, is this. Everyone acknowledges that lots of things, including lots of things outside the head, *cause* cognitive processes. Books, faraway rocks, and the word "Rosebud" can all cause cognitive processes, in fact, can cause certain cognitive processes, rather than others. But not all that many things *constitute* cognitive processes. And certainly not everything that causes cognitive processes constitutes cognitive processes. Indeed, one might even think that being such a cause *precludes being such a constituent*, causes being, after all, distinct from what they cause. (Following this thought suggests, to me at least, a deep confusion in the kind of juxtaposition of cause and constituent in play that I will not pursue here.) Thus, we need a *principle of demarcation* that delineates mere causes from genuine constituents of cognition.

There are two short but basic questions to pose here. First: why? Second, one that I have picked up at occasional postmodernist seances: who are "we"? Take them in reverse order. Precisely who needs this kind of principle of demarcation? Surely not practicing cognitive scientists, in part because they seem to have gotten along perfectly well without one until now. Philosophers? Well, not philosophers who take one of the chief lessons of the failure of logical positivism in the philosophy of science, the collapse of the analytic-synthetic distinction along Quinean lines in the same, and the limitations of conceptual analysis to be a deep suspicion of the search for such principles.

What prompts the felt need for this particular principle of demarcation? Here Rupert is somewhat torn, as are other recent critics of HEC. On the one hand, questions about the "mark of the mental" (p.15) are presented simply as part of the general quest to understand the nature of the mind and cognition. On the other hand, a principle of demarcation distinguishing mere causes from genuine constituents of cognitive processes is required *if one is to take the extended mind thesis seriously as a proposal within the cognitive sciences*. For once we go extended, then there are so many more causes to sort out from those genuine cognitive constituents.

Successful attempts to articulate even *a* mark of the mental systematize large chunks of cognitive phenomena. Big tick for Brentano here in appealing to intentionality, and at least half-marks for those who have more recently posited some relationship to consciousness. Those successes regiment thinking about cognition and the mind, but they do not settle more intricate questions about those phenomena. In particular, just as they do not settle *precisely what kinds of agents have cognition and minds* (e.g., individuals vs groups? biological vs artificial?), they do not determine *the boundaries of the mind and cognition*. Worth keeping in mind here is that since extended cognitive systems are extensions of existing in-the-head cognitive systems, they will inherit whatever mark of the mental those systems have.

Consider what is involved in taking *extended or non-extended* cognition seriously. One might think that it is just obvious that internalists must face a version of the cause-constituent demarcation problem -- not all causes of cognitive processes are outside of the head, after all. In fact, some may even think that this problem has been solved, or close enough, within one or another internalist framework. Maybe. But the solutions that are close enough, or at least on the right track are *not* those that give a one-word answer to questions like "What is the mark of the mental?" Rather, they are those that result from detailed empirical work that identifies *particular cognitive systems inside the head*, of which there are many. Things that fall outside the physical boundaries of those systems are potential causes, and those that fall inside those boundaries are potential constituents, of the specific cognitive processes that those systems generate and facilitate the completion of. Precisely the same is true of extended cognitive systems. But there is no one-size-fits-all answer appealing to mental marks or demarcation by cognitive systems that is both empirically adequate to in-the-head cognitive science and antithetical to HEC.

This then brings me to the second major argument in Part I, which relies heavily on the notion of a cognitive system. This argument is introduced in Chapter 3, and expanded in detail at the beginning of Chapter 4. Despite rejecting the basic premise of this argument -- that we require a systems-based principle of demarcation -- because of scepticism about the need for *any* principle of demarcation here, I am on board with Rupert about the importance of *cognitive systems* for both empirical and conceptual work on cognition. I also concur that cognitive systems are integrated (cohesive) wholes, physically-bounded, reliable, mechanistic in their operation, and typically or paradigmatically possessed by body-bound organisms and perhaps other like individuals. What I don't see is how *any of this* determines where the boundaries of those cognitive systems are, let alone that those boundaries always or even often fall within the physical boundaries of the organisms that possess the corresponding cognitive capacities.

Many, many cognitive systems do, of course, fall within the organismic envelope. Yet this is not entailed, or even made more likely, by the nature of cognitive systems or cognitive processing. Insofar as these have a nature, functionalists and materialists are on the right track, as are the particular kinds of functionalists who are computationalists about cognition. But none of this gets you to an individualistic view of cognitive systems, and so cognition, any more than do appeals to "causal powers" (Fodor 1987: ch.2; Wilson 1995: ch.2) or to computationalism (Wilson 1995: ch.3), or to other general metaphysical or methodological notions in play in the long-standing debate between individualists and externalists.

One characteristic of cognitive systems that Rupert emphasizes about which I have more doubts is *temporal durability*. In the explicit argument that he labels the *Argument from Empirical Success and Methodology, Redux* (pp.59-60), Rupert appeals to this notion by talking of a cognitive system's *temporal grain*, arguing, somewhat indirectly, that extended systems have temporal grains that are too fine-grained: in short, extended systems are not sufficiently durable. The argument begins:

Premise 1. A significant amount of successful research in cognitive psychology (and allied fields) presupposes the existence of human cognitive systems of temporal grain δt .

Premise 2. A causal principle of systems individuation [used by HECers] yields systems of temporal grain smaller than δt .

Premise 3. Individuating cognitive systems such that their grain is smaller than δt compromises the research successes referred to in *Premise 1*.

Here I have two related worries. The first is that the explicit argument itself suggests something like a classic mixed quantifier confusion between the plausible claim that for any cognitive system, there is an appropriate temporal grain or durability range for that system, with the much stronger and quite implausible claim that there is one appropriate temporal grain or durability range for any cognitive system. This may be merely suggestive and so best put aside, but it also draws attention to the second problem: that the plausible

reading of Premise 1 provides the basis for a criticism of extended cognition only if Premise 3 is interpreted in a way that presupposes that extended cognitive systems never or seldom have "the appropriate" temporal durability. If that is true, however, then the argument is question-begging. Rupert takes the claim about the inappropriate durability of extended cognitive systems to rest on the empirical success of traditional cognitive science, but it is at best unclear how this appeal could support what appears to be quite a strong and general claim about the temporal grain of extended systems.

Andy Clark and I have defended the view that durability is one of two primary dimensions in terms of which we can taxonomize extended cognitive systems, arguing that we can have what we called one-off, repeatable, and semi-permanent extended cognitive systems (Wilson and Clark 2009). Rupert clearly rejects one-off extended cognitive systems, and maybe he is correct to do so. Perhaps our labeling here was unnecessarily provocative, for our chief point was twofold: that just how often a cognitive system is realized, or put to work, does not determine whether we have an extended cognitive system rather than merely an external cause plus an internal cognitive system; and that there has been an undue emphasis in the literature on relatively permanent kinds of extended systems, something encouraged by talk of extended *minds*, rather than extended *cognitive systems*. Insofar as Rupert's book shifts the focus to (or perhaps keeps it on) cognitive systems, we head nod. But insofar as it combines that focus with very general metaphysical and methodological appeals in order to undermine HEC, at least I head shake.

The appeal to temporal grain, and at least my second concern about it, thus point to a general difference between my own view of HEC and that of many of its critics, including Rupert, that affects my assessment of much of his book. Key notions in thinking systematically and critically about cognition -- like cognitive system, realization, natural kinds, dynamical systems, functionalism, materialism, consciousness, intentionality, to take concepts that play important roles in *Cognitive Systems* -- warrant philosophical analysis of a kind that is informed by, and in turn informs, the cognitive sciences. Here Rupert and I agree. But I see the chief role of such analysis, including my work on realization that Rupert discusses in detail in Chapter 4, to be to open up possibilities for such empirical work, either by revealing existing constraining assumptions as having dubious epistemic credentials, or by offering sufficiently general construals of these notions to support a pluralistic vision of the cognitive sciences.

Rupert seems to think that the philosophical work to be done at this fairly general level will determine whether substantive views, such as HEC, are *true* or at least *likely to be true*. To his credit, Rupert identifies aspects of explanatory practice in the cognitive sciences relevant to the assessment of the extended cognition thesis, such as the identification of cognitive systems with durability. While no one holds that whether this or that aspect of the cognitive sciences is individualistic or externalist can simply be read off from the explanatory practices in which cognitive scientists engage, the pluralism intrinsic to the version of extended cognition that I embrace sits ill at ease with the project of either establishing or refuting HEC by considering appeals to natural kinds, developmental systems theory, dynamical systems, or consciousness.

One final consideration in play in debates over HEC, one neglected by Rupert, is the parallel between HEC and the hypothesis of extended *biology* (HEB; see Wilson 2005). Although Rupert briefly touches on HEB in several places (e.g., pp.79-80 on fitness; pp.113-118 on developmental systems theory in biology), it is not clear what his global view of HEB is, or whether he has such a view. HEB receives substantial support from emerging work in various areas in the biological sciences, and exemplifies just the sort of pluralism about the biological sciences that we should take more seriously when thinking about the cognitive sciences. Yet HEB has not given rise to philosophical resistance of the kind that partially constitutes the debate over HEC in the cognitive sciences. I suspect that more sustained discussion of HEB and HEC, and more generally of the biological and cognitive sciences together, would have altered not only the kinds of arguments that Rupert provides and considers, but many of the conclusions he argues for with respect to cognitive systems.

Finally, what about the existential angst caused by those damned covers? Since there is nearly always method in madness, I draw attention to the fact that two recent defenses of HEC, Andy Clark's *Supersizing the Mind* (OUP, 2008) and Alva Noë's *Out of Our Heads* (2009) both feature what we might call exploding heads on their covers. Ah, the pattern becomes clearer now: for critics of extended cognition, nothing to do with heads or brains, while HECers are represented by heads that go off firework-like, intimating hives of beyond-the-head activity. A final puzzle for now, given the systematic nature of our study, which I leave as a take home exercise for the reader (with a hat-tip to John Sutton): explain why Dan Dennett's classic *Brainstorms* features three winter-clothed people and a sheep. A Beatles allusion? A wry Biblical reference? A nod to the New Zealand branch of the Australasian Association of Philosophy?^[1]

References

- Adams, Fred, and Kenneth Aizawa, 2008, *The Bounds of Cognition*. Walden, MA: Wiley-Blackwell.
- Clark, Andy, 2008, *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. New York: Oxford UP.
- Fodor, Jerry A., 1987, *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.
- Gallagher, Shaun, and Dan Zahavi, 2008, *The Phenomenological Mind*. New York: Routledge.
- Noë, Alva, 2009, *Out of Our Heads: Why You are not Your Brain and Other Lessons from the Biology of Consciousness*. New York: Hill and Wang.
- Wilson, Robert A., 1995, *Cartesian Psychology and Physical Minds: Individualism and the Sciences of the Mind*. New York: Cambridge UP.
- Wilson, Robert A., 2004, *Boundaries of the Mind: The Individual in the Fragile Sciences: Cognition*. New York: Cambridge UP.
- Wilson, Robert A., 2005, *Genes and the Agents of Life: The Individual in the Fragile Sciences: Biology*. New York: Cambridge UP.
- Wilson, Robert A., and Andy Clark, 2009, "How to Situate Cognition: Letting Nature Take its Course" (with Andy Clark) in Philip Robbins and Murat Aydede (editors), *The Cambridge Handbook of Situated Cognition*. New York: Cambridge UP, pp.55-77.

[\[1\]](#) Thanks to Rob Rupert for feedback on a draft of the review, and John Sutton for the same, as well as for getting me to think more deeply about book covers.