

©2009 The Institute of Mind and Behavior, Inc.
 The Journal of Mind and Behavior
 Autumn 2009, Volume 30, Number 4
 Pages 313–330
 ISSN 0271–0137

Supersizing the Mind: Embodiment, Action, and Cognitive Extension. Andy Clark.
 New York: Oxford University Press, 2008, 286 pages, \$35.00 hardcover.

Reviewed by Robert D. Rupert, University of Colorado, Boulder

For well over two decades, Andy Clark has been gleaning theoretical lessons from the leading edge of cognitive science, applying a combination of empirical savvy and philosophical instinct that few can match. Clark's most recent book, *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*, brilliantly expands his oeuvre. It offers a well-informed and focused survey of research in the burgeoning field of situated cognition, a field that emphasizes the contribution of environmental and non-neural bodily structures to the production of intelligent behavior. The situated research program, fledgling though it may be in some respects, has reached an age at which its philosophical stock can reasonably be taken; and Clark is just the person to take it.

Supersizing the Mind consists of three main divisions. The first develops the case for the distinctively extended view of cognition, according to which the human mind or cognitive system (or human cognitive states or processes) literally comprises elements beyond the boundary of the human organism. The second responds to critics of the extended outlook: Frederick Adams, Kenneth Aizawa, Keith Butler, Brie Gertler, Rick Grush, and me, among others. The third major division evaluates nonextended strands in the situated program, in particular, those that emphasize the role of the non-neural body in cognition.

Clark draws conservative conclusions with regard to some issues. For example, many situated theorists eschew functionalism, computationalism, and talk of mental representation. Clark rejects standard situated criticisms of these theoretical approaches, citing empirical work in response (pp. 19, 24):¹ he repeatedly illustrates the constructive role that the tools of cognitive scientific orthodoxy can play within the situated program. Here I find much of what Clark says illuminating and uncontroversial.

Clark flies radical colors, however, in respect of the extended view. Moderate versions of the nonextended, situated approach — which simply emphasize the subject's interaction with the environment during problem-solving — have guided some fascinating research in cognitive science; moreover, we are quite used to thinking of the body as

Thanks to Bryce Huebner for feedback on an earlier draft of the material in "Getting the Challenges Straight." Requests for reprints should be sent to Robert D. Rupert, Ph.D., Department of Philosophy, University of Colorado, UCB 232, Boulder, Colorado 80309–0232. Email: rupertr@colorado.edu

¹All page references are to *Supersizing the Mind*, unless otherwise noted.

the locus of mind and self, and thus this work does not strongly challenge commonly felt intuitions about the mind's location. In contrast, the extended view faces a steep climb. Few, if any, research successes can be attributed specifically to the extended outlook. Furthermore to many in the philosophical community, the extended view seems incredible on its face. Could my mind — the conscious, emotional, problem-solving self — literally be composed partly of hard drives and cell phones, notebooks and sketch-pads? Why is this not gratuitous mystery mongering? For various theoretical reasons, I think we should remain open to the possibility that human cognition is substantially extended, enough so as to make a difference in cognitive science. Yet, the advocates for the view owe us strong arguments in its support. Clark unapologetically embraces this challenge.

The present essay consists of three main parts. The first section focuses on *Supersizing the Mind's* arguments for the extended view. In the second, I defend myself against Clark's rejoinders. The third summarizes some of the themes and arguments found in *Supersizing the Mind's* final and shortest division.

Action and Interaction: The Case for Genuine Extension

In this section, I present the main themes of *Supersizing the Mind's* first division (chapters 1–4). I sometimes associate these themes with specific chapters, but in reality, most of these ideas appear throughout the first four chapters and, in some cases, in later chapters as well. A further note about strategy and structure is in order. The discussion to come weaves together exposition of Clark's views and critical responses to them, and thus it may be worth bruiting, up front, some of the questions that animate many of my critical concerns: Do actual situated successes support the extended view of the human mind and its cognitive processing? If such results have been produced, are they substantial enough to suggest a revolution (or paradigm shift) in cognitive science? Or, do the successes of the situated program instead show something more modest — that the brain or organismically bound human cognitive system makes much greater use of environmental structures than one might have expected?

Consider the primary theme of chapter 1, that of information self-structuring. Here is one version of the thesis of self-structuring:

The embodied agent is empowered to use active sensing and perceptual coupling in ways that simplify neural problem solving by making the most of environmental opportunities and information freely available in the optic array. (p. 17)

This sort of active sensing comes in a variety of forms, but two aspects of it are central to Clark's presentation: (a) that the cognitive system detects correlations between its self-generated activity — movement — and the resulting perceptual or kinesthetic signals and (b) that the agent can learn effectively by intentionally moving in various ways so as to try to produce data that exhibit such correlations.

Cognitive scientists should pursue, full throttle, research on the self-structuring of input. I see little connection, however, to the extended view. Much human learning is clearly a matter of detecting the kinds of correlation Clark describes, but these correlations hold *between structures within the organism*; in Clark's examples, the events that constitute learning all amount to the recording of correlated patterns of activity within the organism or, in cases of artificial intelligence, within a neatly bounded artificial system (a connectionist network, in one of Clark's examples). Certainly external material

plays a historical role in producing those traces (cf. Rupert, 1998), but this is merely a point about the causal forces affecting the development of the organism or artificial system. Such a historical point is not likely to impress or surprise anyone — even the staunch, internally oriented nativist recognizes the importance of environmental stimulation in intellectual development. Wisely, Clark does not take the extended view to be a thesis about the subject's history of causal interaction with the environment (p. xxvii); however, what, if not historical, is the role of external material as it contributes to learning via informational self-structuring?

Let me cast this concern a bit differently. In Rupert (2004), I distinguished between HEC — the hypothesis of extended cognition — and HEMC — the hypothesis of embedded cognition. The former is the extended view as described above. The latter, HEMC, holds that the human cognitive system is organismically bounded but that it interacts to a surprising extent with external materials in the course of cognitive processing. When reading *Supersizing the Mind's* chapter 1, I repeatedly found myself thinking that Clark had provided clear examples of HEMC-based, but not HEC-based, cognitive processing. Clark makes the point as effectively as anyone. Time and again, Clark describes the relevant processes in terms that clearly favor a HEMC-based approach. Here is Clark, quoting Lungarella and Sporns: “the agent's control architecture (e.g., nervous system) attends to and processes streams of sensory stimulation, and ultimately generates sequences of motor actions which in turn guide the further production and selection of sensory information” (p. 17). The control architecture issues motor commands and, as a result, indirectly produces sensory stimulation — and the commands, the stimulation, and the resulting correlations between them are all internal. Clark goes on to describe research by Fitzpatrick and Arsenio that involves “the cross-modal binding of incoming signals” (p. 18). In what sense are they incoming? In the standard sense: they enter into a robot's computational system through peripheral sensory channels (or are produced internally via proprioception). Over the following pages (pp. 19–21), this theme recurs in a handful of further examples, always to the same effect: the robot or the network or the child learns by interacting with the environment. What does this learning amount to? A rich way of extracting correlations from the incoming data (broadly speaking) and encoding them. Similar remarks apply to the later discussion of the learning of sensorimotor contingencies (p. 23). What is it to learn such contingencies? It is to have the physical materials of one's body, mostly one's brain, altered in certain respects. This is clearly an internalist view, HEMC, not HEC.

Move now to chapter 2, the primary thesis of which concerns the extent to which the boundary of the human body can effectively shift. Clark calls this the “negotiable body” and offers two primary sorts of illustration. In the first category, consider sensory enhancement, the most striking examples of which involve artificial sensory inputs. For instance, a grid of pointed objects can pass patterns of electrical stimulation to the skin, which, with some training, can be used by a subject as surrogate sensory stimulation; if, for instance, the patterns delivered reflect the shifting patterns produced by a head-mounted camera, a blind subject can use the skin-delivered patterns of pulses as visual inputs (pp. 35–36). In some cases, such patterns of electrical stimulation are delivered straight to cortex (p. 36).

The second kind of illustration involves something more like one's representation of one's own bodily boundary and one's control of it. Under certain conditions, the brain incorporates tools into the body schema, that is, represents them as extensions of the subject's own body. For instance, neurons in macaques trained to retrieve food using rakes take on new receptive fields in such ways as to suggest that the trained

macaques' brains treat the rakes as extensions of their own hands (p. 38). Prior to training, certain bimodal neurons are distinctively sensitive both to touch on a particular area of the hand and to visual stimulus of an object approaching that same part of the hand. After training, these neurons are specially sensitive to visual stimulus of objects in the vicinity of the rake head, in the same way they previously had been to visually presented objects near the relevant portion of the hand. The monkeys' brains seem to shift their representation of the body from the hand, as locus of interaction with the environment, to the end of the rake, as the new locus of interaction.

Both kinds of example seem subject to a convincing deflationary diagnosis. In connection with various examples of the first kind, Clark says, "What matters, in each case, is closed-loop signaling so that motor commands affect sensory input" (p. 37). This, however, suggests HEMC, not HEC. Patterned input is delivered by unusual means, and the subject learns to correlate this input with motor signals or other sensory signals. In the second kind of case, it seems equally clear that the interesting story is nonextended — in fact, this seems to follow from the very nature of the evidence at issue. Research on neurons in macaques' intraparietal sulcus may show that macaques represent their bodily boundaries differently after training, but to the extent that the research shows this, it does so by showing *that macaques use neural resources to represent their bodies in a new way*; and neural resources are, of course, inside the organism. Internal, neural resources do the cognizing: they represent bodily boundaries, track ongoing activity of the body, and send motor commands to "body" parts, whether or not the parts so commanded are components of the organism.

Much of the temptation to construe the second kind of case as the basis for an extended view seems to rest on a vehicle-content confusion. The empirical work cited by Clark places the vehicles of learning — the vehicles between which correlations are detected (by the use of further vehicles) — inside the organism, in the brain. Of course, in some cases the *content* carried by such vehicles would, were it accurate, speak in favor of the extended view. The content of a collection of bimodal neurons might be, "the rake is part of my body." That, however, bears little on the actual location of cognition. My neural vehicles might represent my cognition as taking place on the other side of the planet, but thinking doesn't make it so. The processing — that is, the causal interactions among vehicles carrying these contents — is all in my brain, and the very work to which Clark appeals seems to make this clear.²

To be fair, chapter 2 contains intimations of at least two further arguments, and perhaps these connect Clark's empirical examples directly to HEC. The first of these arguments has a phenomenological feel. As Clark observes, "The typical human agent, circa 2008, feels herself to be a bounded physical entity in contact with the world through a variety of standard sensory channels It is a common observation, however, that the use of simple tools can lead to alterations in that local sense of embodiment" (p. 31 — and see pp. 9–10, 33, 37, 80). Here Clark is introducing his discussion of the way in which extended cognitive systems — as new systems in the world — take shape. One might infer from this, then, that Clark embraces an argument of the following (now stripped down) form: it seems to human subjects that their cognition is extended; therefore, human cognition is extended.

²Clark is well aware of the risk of confusing vehicles with their content. He himself points out the difference between the representation of green and the actual color. The representation itself need not be green in order to stand for or mark the color green in our thought processes (p. 57; also see p. 76).

In an endnote, however, Clark (p. 238) disavows the straightforwardly phenomenological argument for the extended view. Given that Clark repeatedly employs the language of phenomenal experience to describe the extended view or to prime readers' intuitions in favor of it, this disavowal might seem an anomalous afterthought. A more coherent and charitable interpretation, however, takes Clark to be developing an inference to the best explanation: subjects have certain phenomenological experiences of their own cognitive activity, and these experiences are best explained by cognition's actually being extended (pp. 41, 81). This style of argument is not made explicit, though; and most importantly, Clark offers no independent, extended account of the production of phenomenological experiences (or intuitions about them) and their relation to neural representations — and certainly no account that is superior to competing nonextended accounts.³ As things stand in *Supersizing the Mind*, then, phenomenological points seem provocative but something of a red herring.

Consider now a (broadly speaking) evolutionary argument. In responding to critics of the negotiable understanding of the body (and perhaps indirectly explaining the phenomenology of negotiability?), Clark draws the reader's attention to the many ways in which internal human cognitive processing is suited to the environment in which it occurs (p. 40). Many cognitive processes succeed only against the backdrop of implicit assumptions concerning the sort of information available from the environment (pp. 40–41, 74–75). How, though, do these observations support the extended view? Clark's most explicit comments in this regard focus on the role of representational resources: "[T]he effect of extended problem-solving practice may often be to install a kind of motor-informational tuning such that repeated calls to epistemic actions become built into the very heart of many of our daily cognitive routines. Such calls do not then depend on . . . *representing* the fact that such and such information is available by such and such a motor act" (p. 75). The idea seems to be that, if a fact about the world is not explicitly represented, yet some cognitive process functions properly only when that fact holds, then the part of the world constituting or described by that fact becomes a literal part of the cognitive process.

This is curious style of argument, resting as it does on one of the central insights of the embedded view: that certain heuristics employed by the local computational (or connectionist, or dynamical) system are valid only when employed in an environment of a certain sort, yet nevertheless are especially useful when employed in that kind of environment (Gigerenzer, 2000; McClamrock, 1995). It seems quite sensible to say that the cognitive system adjusts — either developmentally or evolutionarily — to its environment, but of course this presupposes the existence of a cognitive system that is becoming so suited. To take the tailoring process to bring into existence a further cognitive system serves no purpose. Compare: as one climbs a very high mountain, one's breathing adjusts to the changes in atmospheric pressure and density, but this provides no reason to introduce a new biological unit, the organism-plus-atmospheric-pressure-and-density. Otherwise indispensable theoretical constructs — the organism, its properties, and the ways in which they interact with environmental factors — do all of the explanatory work needed; it seems gratuitous to recognize all of this, then add a new biological system: the organism-plus-atmospheric-pressure-and-density.

³At one point, Clark (p. 81) refers to earlier sections (1.3 and 1.4) in connection with phenomenal experience. The discussion in the earlier sections does not, however, directly address questions about the production of phenomenal experience so much as it identifies the functional contributors to the solution of certain types of problems — e.g., Ballard's blocks-copying task. At the very least, there is a lot of filling in to do here.

My concerns might seem mere niggling absent an account of what makes something cognitive (or part of a genuinely cognitive system). Limitations of space prevent me from going on at length about these issues (see Rupert, 2004, 2008, 2009), but here is a quick sketch of what I think is a promising and principled view concerning the nature of cognition: something is cognitive if and only if it is the state of a cognitive system, where a cognitive system is the persisting collection of mechanisms the integrated functioning of which causally explains, case-by-case, instances of intelligent behavior. Cognitive processing is not simply the activity of whatever causally contributes to the production of intelligent behavior. What separates the cognitive contributors from the noncognitive ones? Insofar as there is any principled and theoretically useful distinction between the cognitive contributors and those merely used by (or interacted with by) the cognitive processes, it is this: the former, the genuinely cognitive processes, are the activities of the fundamental construct of cognitive science, the cognitive architecture (which Margaret Wilson [2002, p. 630] calls the “obligate system”). This, I take it, provides a (contingent, nondogmatic) reason to reject the extended view; the persisting collection of integrated mechanisms is, for most humans most of the time, entirely within the boundaries of the organism. To the extent that this systems-based proposal recognizes cases of extended cognition, these are virtually irrelevant to the philosophical foundations of cognitive science; they are so few, far between, and atypical that they do not reveal “ourselves more truly as creatures of the world” (p. 232, from Clark and Chalmers, 1998) or lead to paradigm shift in cognitive science.

As noted above, Clark is impressed by the extent to which bodily (including neural) systems evolved — either literally or metaphorically — so as to depend on ecological assumptions (pp. 8, 29, 67, 75, 130, 251). Such observations, though, do not favor HEC over HEMC. To the contrary, it is eminently plausible that the bodily bound cognitive system evolves or develops so as to settle on cognitive “shortcuts” that are effective in the environment in which that system typically operates. If the cognitive system is the relatively persisting package of integrated resources that plays a privileged explanatory role with respect to the production of individual behavioral outcomes, then this embedded gloss of evolutionary facts makes perfect sense. Absent a competing, and at least equally plausible, account from Clark regarding what confers cognitive status on a given causal contributor to the production of intelligent behavior, the nonextended interpretation seems not just optional, but superior.

Another theme touched on briefly in chapter 2 is that of transformation: the appearance of “novel properties of the new systemic wholes” (p. 33) at work in extended cognitive processing. Chapter 3 emphasizes this idea to a much greater extent, in particular, with regard to the transformational contribution of external codes (that is, public languages and other systems of external symbols, such as mathematical symbols — pp. 50–53). Clark argues that these material symbols transform human cognition (pp. 50, 57),⁴ conferring upon humans a wide range of capacities distinctive of human intelligence. It is, for example, only by being able to represent our own thoughts that we humans become capable of the higher-order thinking (in the philosopher’s sense of thinking about thinking) at the root of many of our impressive cognitive achievements (p. 58); and on Clark’s view, we become able to represent our own thoughts only because an external code is available.

Fair enough, but this does not seem to support the extended view. Clark’s points appear to contribute only to a historical, causal story about external contributors to cognition — of the sort set aside above. Consider an entirely orthodox, internalist

⁴Also see pp. 63 and 108–109 for transformation-based arguments in other contexts.

account of the effects of external, material symbols on higher-order cognition. A subject learns a new word that expresses a concept or thought. As part of this process, the subject forms a mental representation of it. That mental representation (a symbol in *Mentalese*, if you will) is associated with the thought expressed by the mental representation's public language counterpart (i.e., the actual word), and it also enters the domain of various computational processes in the cognitive system. As a result of the latter change, the internal, computational cognitive system can compute functions over this new mental representation and thus can indirectly compute functions that take a thought (i.e., something with the content of a first-order thought) as argument; in particular, it can compute functions that construct compound symbolic strings comprising a mental representation of the subject him or herself, of the general activity of thinking, and of the content of the mental state in question. Thus, the human cognitive system can think about its own thoughts. (I haven't explained why it's easier for the cognitive system to manipulate mental representations of words than it is for the cognitive system to manipulate representations of its own thought content that are not patterned after an external code; but neither has Clark, at least if one is concerned with off-line cognitive processing, i.e., processing that takes place in the absence of the external symbolic tokens themselves — see below.)

What is Clark's objection to this relatively mundane, internalist view? I found Clark's discussion in this regard to be uncharacteristically opaque. Clark asserts that, in the case of number words, "there is (at least) an internal representation of the numeral, of the word form, and of the phonetics" (p. 52). This, however, recognizes the essential representational materials posited by the internalist story, which takes the transformational power of language to be merely causal and historical. Psycholinguistics constitutes one of the most influential cognitivist programs, going back to Chomsky's work in the 1950s; and standard psycholinguistic proposals presuppose internal representations of words, including both graphic and phonemic properties of them, as participants in language processing.

What, then, motivates the extended gloss of what otherwise appear to be merely historical and causal processes? Clark's primary objections to the internalist story seem to be that internal representations of words are "shallow, imagistic inner encodings" (p. 238; cf. p. 53) and not, individually, "*fully content-providing*" (p. 52). All of that, however, seems fine from an internalist standpoint (although some might reasonably wonder what being fully content-providing amounts to; is Clark assuming — dubiously, I would say — that every genuine *Mentalese* symbol must enter into all of the internal relations that the subject associates with the external content carried by that symbol?). Internalist architectures and algorithms (in the case of computational architectures) vary quite a bit, so there may be no canonical model of the internal representation of material symbols. Nevertheless, standard cognitivist approaches provide many resources for cognitive modeling. Consider, for example, that computational models commonly incorporate pointers (Newell and Simon, 1976/1997), which seem about as shallow as mental representations get; thus, the shallowness of mental representations of external symbols does not conflict with orthodox approaches in cognitive science. Neither does the imagistic nature of representations of public symbols. Computational primitives need not take any particular form, so long as they're treated as primitives by the computational system. Thus, there is no reason a computational primitive cannot possess pictorial or imagistic properties. Orthodox computational approaches can be indifferent about this issue. So long as the imagistic properties play no role in cognitive processing, then a computational account of that process remains as viable as ever.

But, what if the particular form — the physical implementation or realizer — of a given mental representation (individuated in terms of its content) varies from subject

to subject (say, from the speaker of one language to the next)? That is, what if two subjects form substantially different shallow, imagistic representations of number words with the same content (both referring, for instance, to ninety-eight)? Won't the imagistic features of the representations govern the subjects' responses in at least some circumstances? Perhaps, but that shows only that computationalism leaves something out, not that there is anything extended about the story. It is one thing to say that certain behavioral variables are distinctively affected by a vehicle's physical properties; it is quite another to hold that the vehicle itself is external. In the standard language-based case, the vehicle with imagistic properties is still an *internal* vehicle; nothing Clark says in chapter 3 suggests otherwise.

This being said, I may have neglected an important aspect of Clark's view. Clark is impressed by the way in which external symbols can, when immediately present, seem to play an active, attention-directing role in cognition (pp. 48, 57). All right, but if words do play such a role, they do it via activating internal representations. Clark draws a recurring example from the work of Dana Ballard and his associates (Ballard, Hayhoe, Pook, and Rao, 1997). Subjects are shown a pattern of colored blocks — the target — and are given various further colored blocks as resources to use to replicate the target. In these experiments, the subject is faced with three distinct compartments or areas: one displays an already assembled pattern of blocks; a second provides the subject with a jumble of blocks, resources to be used by the subject to try to replicate the pattern in the first area; and a third begins empty but is meant to be used by the subject to build her replication of the target model. By recording fixation points of subjects' eyes, Ballard et al. showed that subjects often (but nothing close to exclusively) use a strategy that relies more on looking back and forth than it does on the committing of lots of information about the target to internal memory: instead of memorizing, upon one fixation, the location and color of a block in the target model, subjects often look back and forth between the target and the resource areas, apparently holding only the color in memory at the first stage, then looking back to the target to fix a location in memory before placing the block in the workspace.

We should not, however, misinterpret the results. The experiments hardly show that subjects don't rely on mental representations of block colors or positions. To the contrary, one of the commonly used strategies (the P-D strategy — Ballard et al. 1997, p. 732) relies heavily on internal memory. Moreover, this strategy tends to be used more frequently in later parts of each trial, after the subject has encountered the target repeatedly. Thus, it appears that throughout a given trial, subjects build up a persisting representation of the target structure. (Other commonly used strategies, M-P-D and P-M-D, also rely more heavily on internal memory than the strategy that impresses Clark.) Most importantly, even on the least memory-intensive strategy — the one that involves the most looking back and forth — the deictic pointers used by subjects must represent the colors of the external blocks or their position — even if only one block and one property at a time. What's interesting about the pointers is the dynamic reassignment of them to the job of representing various external things: positions, or colors. Each time a visual pointer is "reassigned," however, it has to be bound to persisting representations of properties, or else it is useless in the copying task. Comparing two bare pointers to each other or comparing one bare pointer (aimed at the color of a block the subject has just attended to) to the color of a block in the resource pool does not do the subject any good. The subject has to be able to "decide" whether the pointer and the visual representation of the color of the block to which she is currently attending (while looking at a candidate block in the resource pool) are the same — so that she can pick up the correct block. This requires binding the pointer to an external object but also to an internal representa-

tion of its color. Ballard et al. do not deny this; in fact, it's built into their approach (p. 725).

Return now to the case of words. When reading, some words differentially capture the subject's attention. Nevertheless, it's reasonably clear that mental representations of words commonly contribute to cognitive processing in the absence of the actual units; during literature exams, students routinely produce names of characters and descriptions of settings, without having the text at hand. So, there is independent reason to posit mental representations built up by subjects while reading. In which case, the attention-directing role of external resources begins to look pretty humdrum: when one looks at a given word, it "directs one's attention" by causing the activation of an internal representation of that word. This is a standard internalist view of the causal role of external codes.

Clark intends that chapter 4 make the case in favor of the extended view. Yet, it is largely a compendium of the points made in chapters 1 to 3, so I won't rehash what I've said about those earlier chapters. It should suffice to point out that, to the extent that I have raised legitimate concerns about an extended interpretation of the material presented in chapters 1 to 3, I have called into question the extended conclusion reached in chapter 4.

Getting the Challenges Straight

In the second of *Supersizing the Mind's* three major divisions, Clark responds to critics. Chapter 6, in particular, provides a sustained rejoinder to some of my concerns about the competition between HEMC and HEC. In this section, I evaluate those rejoinders. Such measure can be taken, however, only after first getting clear about the argument to which Clark is responding.

Here, then, is the original argument, in outline:

Premise 1. HEMC is a competitor to HEC (as a philosophy of cognitive science, or as a principled treatment of actual human cognition) [Rupert, 2004, pp. 395–397].

Premise 2. Conservatism favors HEMC (*ibid.*, pp. 395, 405–406).

Intermediate conclusion 1. Thus, other things being equal, HEMC should be favored over HEC (*ibid.*, p. 395).

Intermediate conclusion 2. Thus, we should accept Clark and Chalmers's (1998) endorsement of HEC only if Clark and Chalmers provide a substantial argument in support of HEC (if they can show that other things are not equal).

Premise 3. Clark and Chalmers offer two identifiable arguments in support of HEC, one intuitive and commonsensical (the Otto's notebook argument and the intuitive criteria derived from it), the other an argument in philosophy of cognitive science (the natural- or causal-explanatory-kinds argument) [*ibid.*, pp. 401–407].

Premise 4. The intuitive argument falls to a dilemma: either one accepts Clark and Chalmers's (1998) conscious-endorsement criterion, effectively collapsing HEC into HEMC, or one rejects the conscious-endorsement criterion leaving in place only the remaining three criteria, which are subject to obvious counterexamples (*ibid.*, pp. 402–405).

Premise 5. The natural-kinds argument falls to a dilemma: either we individuate the relevant causal-explanatory kinds in a fine-grained way, in which case Clark and

Chalmers's argument falls to the worries about fine-grained differences, or we treat the relevant causal-explanatory kinds as coarse-grained — which I called “generic” — kinds, thereby robbing such kinds of any causal-explanatory power not accommodated by HEMC (*ibid.*, sections V to VIII; for these two options explicitly presented side-by-side, see pp. 407, 418–419, 424).

Conclusion. Therefore, we have no good reason to accept HEC, given the current state of the evidence (*ibid.*, p. 428).⁵

So far as I can tell, the original paper exhibits this structure clearly. Nevertheless, much of the critical reaction in the literature—including Clark's in *Supersizing the Mind* — rests on a misunderstanding of the role played by sundry components of the argument. I am variously charged with misinterpreting the Parity Principle (more on which below), believing that HEC requires fine-grained similarities between the inner and the outer (in some cases, because I have allegedly misunderstood the Parity Principle), privileging some kind of inner Cartesian Theatre, or obsessing over the organismic boundary. All of these charges rest on a confused — in the second and most common case, an incoherent — reading of the original paper.

Consider the Parity Principle. In “The Extended Mind” (reprinted as an appendix in *Supersizing the Mind*), Clark and Chalmers make the following claim:

If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process. (p. 222)

Clark accuses me of resting some of my concerns on a “persistent misreading of the parity claim” (p. 114). In developing Premise 5's dilemma, I focused at some length on significant dissimilarities between internally realized cognitive states and those external states purported to be cognitive; and as Clark sees things, this focus on fine-grained dissimilarities arises because I mistakenly take the Parity Principle to entail a fine-grained similarity between internal and external cognitive states (pp. 112–115; for similar interpretations, see Menary, 2006, pp. 333, 339–340; Rowlands, 2009, p. 3; Wheeler, *in press*, pp. 3–4).

This diagnosis misses the mark badly. First off, notice that, even though I used a fairly heavy style of citation, I did not once cite or quote Clark and Chalmers's passage expressing the Parity Principle, nor did I paraphrase it. More to the point, consider what I did cite. I explicitly set out to criticize Clark and Chalmers's (1998, pp. 13–14) argument from natural — or causal-explanatory — kinds, quoting at length their statement of the argument, then clearly paraphrasing it (Rupert, 2004, pp. 406–407); it is difficult to see how a reader could have proceeded through the lengthy discussion that follows, with its repeated reference to natural kinds, without having understood this framing of the issues. According to the distinctive premise of Clark and Chalmers's argument, cognitive science benefits from construing its causal-explanatory kinds in such a way that many external and internal states are of the same natural kind or instantiate the same theoretically relevant cognitive property. In response, I

⁵In its closing pages, Rupert (2004) also suggests a more direct argument against HEC, an argument appealing to the privileged role of persisting systems in cognitive science. This concern has been developed in more detail in Rupert (2008, 2009). Below I discuss Clark's response to this argument.

offered Clark and Chalmers a choice: characterize the relevant causal-explanatory kinds in terms of fine-grained functional properties (in terms of the reaction times they support, for instance) or instead opt for a coarse-grained conception of cognitive kinds (the generic kinds of Rupert, 2004, section VII); these two options are explicitly contrasted at least three times (Rupert, 2004, pp. 407, 418–419, and 424). I argued that neither alternative offers both (a) extended kinds (i.e., kinds that singly subsume both internally and externally realized states in a significant number of actual human cases) and (b) a resulting causal-explanatory advantage; the first alternative is not likely to yield extended natural kinds at all (given the sorts of fine-grained properties of interest to cognitive psychologists), and the second alternative yields extended kinds unlikely to do substantive causal-explanatory work. This debate about natural kinds (or properties) has nothing directly to do with the Parity Principle and, *a fortiori*, nothing much to do with a misreading of it.

All of this speaks against the charge, even more common than anything to do with the Parity Principle, that I attacked HEC for a commitment to similarity in fine-grained properties among internal and external states of the same type (Bartlett, 2008, p. 171; *Supersizing*, pp. 112–115; Levy, 2007, pp. 58–59; Menary, 2006, pp. 339–340; Sprevak, 2009, pp. 506–507). Nothing in my discussion presupposes that HEC entails fine-grained similarity of internal and external states. I give no argument of the form “if HEC, then fine-grained similarity; no fine-grained similarity; therefore, not HEC.” Interpretations stating otherwise make a hash of the structure of sections V–VIII of Rupert (2004), most obviously the discussion of generic kinds. Why would I have considered the second horn of the dilemma, the possibility that Clark and Chalmers have in mind generic kinds, if I were taking HEC to entail fine-grained similarity of inner and outer kinds? After all, generic kinds do not require fine-grained similarity. If I had thought HEC entails fine-grained similarities, my treatment of generic kinds would have been much different: I would simply have pointed out the “incoherence” of arguing for HEC by appeal to generic kinds. I did not, however. Rather, I argued that generic cognitive kinds fail to play a substantive causal-explanatory role in cognitive science, exactly in keeping with a consideration of Clark and Chalmers’s natural-kinds argument. It is most uncharitable, then, to interpret Rupert (2004) as resting its criticism of the natural-kinds argument for HEC (or a criticism of HEC itself) on the assumption that HEC entails fine-grained similarities.⁶

Why the widespread misinterpretation? Perhaps it was an oversight on my part not to have discussed the Parity Principle. In connection with the dialectical role of the Parity Principle, there may be a gap between the assumptions I take on and those of (some of) my critics — a gap that naturally leads critics to think I *must* be talking about the Parity Principle, at some level, even if I have not indicated in any way that I am doing so. Methodologically, I am inclined toward a fairly extreme naturalism in philosophy, in keeping with which one would expect both that (a) the strongest arguments for HEC arise from the scientific work and its attendant taxonomy and (b) very little argumentative mileage can be gotten out of the Parity Principle, that is, out of our intuitions about what we would consider cognitive if it were in the head. After all, it’s not up to our intuitive judgments to decide what cognition is; the property *being cognitive* is a scientific construct, vindicated only by the causal-explanatory work it does. In contrast, many philosophers have been trained to take reflective (and even

⁶A further exegetical error is to be found in this vicinity. My discussion of natural kinds does not present an argument against HEC; its immediate role in the paper is to challenge a premise of an argument for HEC, and the criticism of an argument for P is not an argument against P itself.

hip-pocket) intuitions quite seriously; so they might tend to be impressed by the Parity Principle in a way that I am not. As I read it, the Parity Principle by itself is uninformative: if we don't already know which properties determine something's cognitive status, then we do not know whether the intuition that an external process, were it in the head, would be cognitive is evidence for the process's being cognitive when it's external. We don't know whether the intuition "would be judged cognitive if it were in the head" is driven by mere change in location — in which case, the intuition is likely to be a reliable guide to the cognitive status of something external — or is driven instead by our sensitivity to a correlation between putting something in the head and that thing's thereby acquiring properties that are distinctively cognitive (e.g., becoming integrated into a relatively persisting cognitive system); if the latter consideration drives our positive intuition — viz. "yes, it would be judged cognitive if it were in the head" — this intuition itself does not bear on the question of whether the process is cognitive when it is external.

Alternatively, it might be that Adams and Aizawa's (2001) original paper, which appeared before mine, colored readers' interpretation of mine. Adams and Aizawa appealed to fine-grained differences between inner and outer states in what seemed like a direct attack on HEC; and this might have led readers to think I was presenting only a more elaborate version of their argument. Here I simply encourage interested parties to take in hand the outline, given above, and re-read Rupert (2004) comparing the outline with the text.

Admittedly Rupert (2004) is long and filled with asides. Some of these are of interest in the present context, some are not. For example, I attacked epistemic-dependence arguments (*ibid.*, p. 396), claiming that one cannot effectively support HEC by observing that, in order to understand cognition, one must understand the organism's interaction with the world. In *Supersizing the Mind*, Clark continues to advance epistemic dependence arguments (pp. 116, 157–158), for reasons that are unclear to me. I also distinguished HEC from semantic externalism, as Clark and Chalmers do, partly with the intention of showing that even if one accepts HEC, one might get an externalist semantics for the external vehicles that participate in cognitive processing. Moreover, I recast some of Clark and Chalmers's arguments in explicitly functionalist terms, arguing that this tack bears no fruit. All of this may have been a distraction to readers trying to map out the structure of the paper.

Let us turn now to *Supersizing the Mind*'s direct responses to my systems-based view (see footnote 5; above). The emphasis on systems is meant to provide independent grounds for a HEMC-based, as opposed to a HEC-based, approach. As Clark sees things, however, the HEMC-based approach elevates "anatomic and metabolic boundaries into make-or-break cognitive ones" (p. 138); but of course it does no such thing, at least not if "make-or-break" implies that the barrier is absolute or that some interest in the barrier itself drives the arguments in favor of HEMC. The arguments for the systems-based approach (Rupert, 2009) and thus, indirectly, for HEMC rest on (1) the privileged role of the persisting integrated architecture, (2) longstanding and successful uses of the idea of a persisting architecture that interacts with various resources in its environment, and (3) the superfluous nature of a HEC-based redescription of this research strategy. These arguments do not rest on a prejudiced elevation of the boundary of skin and skull. Quite the contrary: they arrive at the nonextended conclusion as a contingent result of past successes and application of methodological principles. The nonextended conclusion follows from negotiable, contingent facts.

Notice, too, that my view depends in no way on there being a Cartesian Theatre or anything like it, in contrast to Clark's suggestion that HEMC depicts "outer

resources as doing their work only by parading structure and information in front of some thoughtful inner overseer" (p. 137). In "The Extended Mind" (p. 17), Clark and Chalmers tentatively suggest that internal consciousness must validate the cognitive status of external states. In Rupert (2004, pp. 404–405), I argued that such a view runs toward HEMC more than it does HEC. We must, however, keep the logic straight here. It is one thing to assume, as I did, that if there is a privileged internal consciousness before which structure and information must be paraded in order that they be cognitive, then HEMC (most likely) wins the day. It is quite another to assume the converse conditional: if HEMC is true, there is a privileged internal consciousness before which structure and information must be paraded in order that they be cognitive. My criticisms of Clark and Chalmers's argument for HEC in no way presuppose the second conditional, which I take to be false. Thus, it is misguided for Clark to level charges of a "magic dust" (p. 136) error in connection with HEMC. The embedded view is consistent with a distributed cognitive architecture; as I have developed the view, HEMC requires only that the architecture, distributed or not, be inside the organism.

Lastly, consider a pragmatic point. Clark sometimes suggests that the adoption of anything short of HEC obscures the importance of the environment from cognitive-scientific view (p. 136). There is, though, no reason to think HEMC occludes the environment's contribution to human cognition. After all, HEMC's expressed agenda is that cognitive science focus on ways in which the human cognitive system interacts with and exploits external resources; it will be an odd HEMC-theorist who ignores the role of the environment. Clark's concern would be more compelling were there actual cases in which the HEC-based perspective led to cognitive-scientific advances and where HEMC, had it been adopted in place of HEC, would have prevented these advances. So far as I can tell, though, the empirical research taken to support HEC was motivated not by a specific commitment to HEC *or* to HEMC, but rather by a general sense that interaction with the environment plays an important role in cognitive processing. Consider the way Ballard and colleagues describe their project in one of the most influential empirical papers in the situated tradition: "Our central thesis is that intelligence has to relate to interactions with the physical world, meaning that the particular form of the human body is a vital constraint in delimiting many aspects of intelligent behavior" (Ballard et al. 1997, p. 723).⁷ This thesis entails neither HEC nor HEMC. Thus, I find no reason to think that, if, on the basis of independent arguments, we adopt HEMC instead of HEC, the context of discovery in cognitive science will be impoverished; HEMC leaves in place the emphasis on interactive processing — the primary theoretical vision driving the research claimed to support HEC directly.

Returning now to the questions about natural-kinds, Clark (p. 115) does ultimately address my central concern — that is, that individuating psychology's kinds in an extended way adds no causal-explanatory power to orthodox cognitive psychology. Clark's discussion in this regard seems best understood as a challenge to Premise 5. We can take him either to be arguing that Premise 5 presents a false dilemma (there is a third way to make extended kinds relevant) or, if the fine-grained v. generic distinction is exhaustive, to be arguing that our individuation of cognitive kinds in one of those two ways can lead to useful HEC-friendly scientific advances.

Clark's parade argument here draws on the work of Wayne Gray and his collaborators (pp. 118–122). In a series of experiments, Gray and associates (Gray, Sims, Fu, and

⁷Cf. the discussion in Rupert (2004, pp. 393–394; footnote 9) of the ambiguity of message in Kevin O'Regan's work on vision.

Schoelles, 2006) measure subjects' tendency to "choose" between the use of internal memory and the accessing of information encoded in external structures. Gray does so by manipulating the relative time-cost of the use of internally and externally encoded information. The results manifest a regular relation: increase the cost of access to environmentally encoded information, and subjects are more inclined to use internally encoded information, and vice versa. The cognitive system seems to "care" about only the time-cost of access to information, not about its location per se. Clark takes this to show that the external locations are part of the cognitive system.

Why, though, should we not take Gray's results to show that, when there is no great cost in terms of time, the cognitive system uses resources beyond its boundary? This interpretation of the results seems equally plausible. Clark seems to need the following premise: a system that uses resources beyond its boundary must (or at least is very likely to) treat the external nature of the location of those resources as intrinsically relevant to the decision whether to use those resources. Otherwise, why would it matter that the system Gray discusses *doesn't* treat this difference as of intrinsic import?

Consider the apparent form of Clark's argument:

Premise 1. If the human cognitive system uses information from two locations, *without treating the difference in locations themselves as relevant*, then both locations are inside the system.

Premise 2. In the cases investigated by Gray and associates, the human cognitive system uses information from two locations, *without treating the difference in locations themselves as relevant*.

Conclusion. Therefore, HEC.

My concern is with *Premise 1*. Why does the failure to treat the two locations as of importantly different kinds, and to instead treat them differently only as a function of time-cost, show that both locations are inside the cognitive system? This seems unjustified. Take a system with any boundaries you like. There almost certainly are cases in which the system accesses information beyond its boundaries and does not treat this difference as anything more than a difference in accessibility (or time to completion, or amount of pain caused in the body to get the information, or whatever). The fact that the system fails to treat the external nature of the information as intrinsically relevant to its decisions shows, so far as I can tell, absolutely nothing about the boundary of the system; it does not show that what we might have thought was an external location is really an internal one. Clark's argument goes through, then, only when *Premise 1* is supported by the implausible assumption that any cognitive system that uses genuinely external information must mark the external nature of the location explicitly and treat it as intrinsically relevant.

Consider, too, the mechanisms by which the cognitive system gains access to information from the internal and external locations. Use of the internal store need not involve the running of on-line perceptuo-motor routines, whereas external stores are accessed only via such routines. With this in mind, it might seem superfluous to characterize the system as accessing an organismically external location, at least if one means to characterize the operative algorithm; if this is correct, the argument just attributed to Clark does not even get off the ground. I see no reason to doubt that, when the system makes use of externally encoded information, the body-bounded system forms internal representations of that information (compare this to the remarks made above about the role of representation in Ballard's results). The process of accessing

an organismically external store, then, may best be cast as an entirely internal process. The system draws on its implicit knowledge of reliable sensorimotor contingencies: motor-command p is followed by a sensory experience that should be treated as the answer to a given query.

Represented in this way, the situation can be described in wholly HEMC-based terms: the use of one internal store contrasts with the use of a distinct internal store. On this way of modeling the computational process in question, the system “chooses” between the retrieval of information from various internal registers: memory register A — part of, say, short-term, declarative memory — and register B — a visual buffer. In cases where B is chosen, the central controller “cares” only that the information shows up reliably in the sensory register, not where the information is in the external world. On this view, both of the locations from which information is accessed are inside the organism, and the system’s process of “choosing” between them has no bearing on HEC.

I am not claiming that the HEMC-based approach is clearly superior, only that it is available. No party to the debate thinks that Gray’s subjects magically gain access to information from the environment. If, however, such access runs via an internal store, then proponents of HEMC will focus on it, rather than the organismically external location of the information — the latter being used in experiments as a way of creating information-bearing states in the organismically internal visual buffer. On this view, the correct representation of the algorithm used by Gray’s subjects includes two locations from which information can be retrieved, both internal. The system can choose between sending motor commands (to, for example, orient the head in a certain way) and waiting for the relevant information to appear in the visual buffer and a second alternative, calling up information from working or long-term memory. Gray’s results show that the time-cost of these two processes can tip the balance in favor of the use of one rather than the other.

The preceding discussion of internal stores, long-term memory, visual buffers, and the like might seem too intellectualist or cognitivist to some readers. Nevertheless, my first concern seems decisive: even if we think in terms of a choice between internal and organismically external locations, this does not increase the likelihood of the extended, relative to the nonextended, hypothesis. Each hypothesis predicts that aspects of human cognition can be effectively modeled using an impartial algorithm. There is nothing in Gray et al.’s work to decide the issue. There is no reason to think that cognitive systems should not treat genuinely external resources impartially; thus, showing that the human cognitive system treats some organismically external locations impartially does not show — or even suggest — that the human cognitive system extends beyond the boundary of the human organism. As I see things, then, the question is whether the external store is part of the persisting set of integrated mechanisms used in overlapping subsets to produce a variety of forms of intelligent behavior (Rupert, 2009).

Embodiment

In this final section, I say a few words about *Supersizing the Mind*’s third main division. This division consists of two substantive chapters (8 and 9), one addressing the sensorimotor theory of perceptual experience (associated with the work of Hurley, Noë, and O’Regan) and one addressing questions about the embodied research program more generally.

The sensorimotor theory of perceptual experience identifies sensory experience with implicit knowledge of sensorimotor contingences — that is, implicit knowledge of corre-

lations between physical movements (or motor commands) and sensations. What one experiences visually, for example, is constituted by one's implicit grasp of how visual stimulation varies systematically with one's movements (broadly construed so as to include, for example, eye movements and muscular commands to hold still).⁸

Clark rejects the sensorimotor view for two related reasons. First, Clark develops a charge of chauvinism (pp. 177–180). As described by the view's advocates, the dependencies in question are a fine-grained affair, and thus the sensorimotor view seems to allow two individuals to have the same type of perceptual experience only if they have highly similar bodies and sensory apparatuses. This concern leaves various responses available to the sensorimotor theorist, bullet-biting among them.

Clark's second criticism, however, seems to decide the case against the sensorimotor theory. Here Clark develops a point found in some of his earlier work (Clark, 2001; also see Block, 2005). It appears that the parts of the visual system related directly to motor control are not the parts of the visual system associated directly with perceptual experience. The perceptual experience of, say, an apple *as such* (i.e., categorized as an apple in way that facilitates the paradigmatic conscious processes of explicit reasoning, planning, and speech production) occurs primarily in the ventral stream, rather than the motor-controlling dorsal stream of visual processing (p. 192). When it comes to visual experience, then, the sensorimotor contingencies as represented in the dorsal stream do not play a constitutive role.

I have one worry about Clark's tack. While convincing on its own, the reader should wonder whether what's good for the goose is good for the gander. At one point, Clark (pp. 191–192) acknowledges a close causal relationship — one of fine tuning — between activity in the dorsal stream and appropriate activity in the ventral stream. He denies, though, that this establishes a constitutive role, with respect to visual experience, for activity in the dorsal stream. Upon reading this, I couldn't help but ask where Clark thinks we should draw the line between merely intimate causal interaction and the kind of dense causal interaction that is supposed to ground constitutive claims. After all, many of Clark's points in favor of the extended view (see, e.g., his summary list on p. 81) emphasize various kinds of close causal interaction. On his view, such patterns of interaction ground constitutive claims regarding cognitive processes (with the consequence that many cognitive processes are partly constituted by external materials); but by what principle does Clark distinguish his favored cases from the sensorimotor theorist's? I'm not claiming there is no such principle, only that readers would be better served were Clark to make that principle explicit and apply it consistently across cases.

In the final substantive chapter, Clark takes on the radical wing of the embodied cognition movement (Shapiro's views, in particular). Philosophers and cognitive scientists who tout the embodied approach often sell it as a solution to the supposed functionalist–computationalist malaise. According to functionalism (and its close cousin computationalism), the body does not matter to cognition, and the study of the mind can proceed entirely independent of the study of the body — or so embodiment's boosters claim. As Clark points out, however, there is nothing in the functionalist program that precludes an interest in the body: "Cognition, for the machine functionalist, was independent of its physical medium in the sense that *if* you could get the right set of abstract organizational features in place . . . you would get the cognitive properties

⁸The sensorimotor view is sometimes expressed in terms of correlations between movements and perceptual experiences, but that, of course, takes perceptual experience to be something independently given, and the resulting view either begs the question or is viciously circular (cf. p. 179).

'for free'" (p. 198). But, of course, you have to have the right sort of material to get those abstract organizational features in place, and finding material that can be so organized may be no trivial matter. Moreover, which algorithms a given mind employs will, on the standard functionalist picture, depend on which kinds of material happen to be at its disposal. As Clark puts it, "[T]he brain's algorithms factor in the bodily structures and opportunities" (p. 204).

Where, then, is the tension between functionalism and the embodied research program? If there is such a tension, it concerns specific patterns of cognitive performance or specific kinds of conscious experience. The proponent of the embodied view might claim that a given cognitive or phenomenal profile *cannot* appear in connection with more than one kind of body. Clark sees this as a kind of mysterianism, though. For, it would seem that virtually any mechanism that a particular kind of body uses to solve a given problem could be mimicked by a different body in which the relevant mechanisms are rearranged physically and temporally (p. 204). Furthermore, with regard to questions about consciousness, Clark reminds the reader of his discussion of the preceding chapter. There are good reasons to think human visual consciousness filters out the sorts of fine-grained bodily features reflected in the sensorimotor contingencies (the features that might tie visual consciousness to one's specific sort of body).

In *Supersizing the Mind*, Clark furthers the debate marvelously. He manages to engage, in a detailed and focused way, with the most active and prominent philosophers writing about situated cognition, while also introducing the empirical work and big-picture ideas that animate the subject. I *highly* recommend this book.

References

- Adams, F., and Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology*, 14, 43–64.
- Ballard, D., Hayhoe, M., Pook, P., and Rao, R. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, 20, 723–742.
- Bartlett, G. (2008). Whither internalism? How internalists should respond to the extended mind hypothesis. *Metaphilosophy*, 39, 163–184.
- Block, N. (2005). Review of *Action in Perception* by Alva Noë. *Journal of Philosophy*, 102, 259–272.
- Clark, A. (2001). *Mindware: An introduction to the philosophy of cognitive science*. Oxford: Oxford University Press.
- Clark, A., and Chalmers, D. (1998). The extended mind. *Analysis*, 58, 7–19.
- Gigerenzer, G. (2000). *Adaptive thinking*. Oxford: Oxford University Press.
- Gray, W., Sims, C., Fu, W., and Schoelles, M. (2006). The soft constraints hypothesis: A rational analysis approach to resource allocation for interactive behavior. *Psychological Review*, 113, 461–482.
- Levy, N. (2007). *Neuroethics: Challenges for the 21st century*. Cambridge: Cambridge University Press.
- McClamrock, R. (1995). *Existential cognition: Computational minds in the world*. Chicago: University of Chicago Press.
- Menary, R. (2006). Attacking the bounds of cognition. *Philosophical Psychology*, 19, 329–344.
- Newell, A., and Simon, H. (1977). Computer science as empirical inquiry: Symbols and search. In J. Haugeland (Ed.), *Mind design II: Philosophy, psychology, and artificial intelligence* (pp. 81–110). Cambridge, Massachusetts: MIT Press. [Reprinted from the *Communication of the Association for Computing Machinery* 19 (March 1976): 113–126.]
- Rowlands, M. (2009). Extended cognition and the mark of the cognitive. *Philosophical Psychology*, 22, 1–19.
- Rupert, R. (1998). On the relationship between naturalistic semantics and individuation criteria for terms in a language of thought. *Synthese*, 117, 95–131.
- Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, 101, 389–428.

- Rupert, R. (2008). Innateness and the situated mind. In P. Robbins and M. Aydede (Eds.), *Cambridge handbook of situated cognition* (pp. 96–116). Cambridge: Cambridge University Press.
- Rupert, R. (2009). *Cognitive systems and the extended mind*. New York: Oxford University Press.
- Sprevak, M. (2009). Extended cognition and functionalism. *Journal of Philosophy*, 106, 503–527.
- Wheeler, M. (in press). Minds, things, and materiality. In C. Renfrew and L. Malafouris (Eds.), *The cognitive life of things: Recasting the boundaries of the mind*. Cambridge: McDonald Institute for Archaeological Research Publications.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin and Review*, 9, 625–636.