The Journal of Mind and Behavior
Summer and Autumn 2010, Volume 31, Numbers 3 and 4
Pages 253–264
ISSN 0271–0137

Cognitive Systems and the Extended Mind. Robert D. Rupert. New York: Oxford University Press, 2009, 288 pages, $55.00 hardcover.

Reviewed by Colin Klein, University of Illinois at Chicago

Robert Rupert is well-known as a vigorous opponent of the hypothesis of extended cognition (HEC). His *Cognitive Systems and the Extended Mind* is a first-rate development of his "systems-based" approach to demarcating the mind. The results are impressive. Rupert's account brings much-needed clarity to the often-frustrating debate over HEC: much more than just an attack on HEC, he gives a compelling picture of why the debate matters.

Rupert's book has three parts. The first part defends his distinctive approach to demarcating the mental. The second part surveys many contemporary defenses of HEC, using Rupert's demarcation criterion to show that they fail to be more satisfying than the *hypothesis of embedded cognition* (HEMC). The hypothesis of embedded cognition, on Rupert's account, situates the cognitive within the boundaries of the organism, but is more sensitive than traditional approaches to the relationship between the cognitive and the external environment. The third part concludes with an elaboration and defense of HEMC, showing that while HEMC is an important advance, it is compatible with computationalism and many of the other theoretical tools offered by traditional cognitive science.

Demarcation and the Systems-Based Approach

In chapter 2, Rupert introduces the notion of a *demarcation criterion* that separates the cognitive from the non-cognitive. Most debates over HEC presuppose such a criterion; most also presuppose that the criterion should come from our best scientific practice. That said, scientific practice is a wooly beast: part of the job of the philosopher is to figure out just what makes good practices good. Here, I think, Rupert is wonderfully explicit. He suggests three *desiderata*, the most important of which is the third: a proposed demarcation criterion should have some empirical payoff: that is, it

---

Thanks to Robert Rupert and Kenneth Aizawa for helpful discussions on an earlier draft. Requests for reprints should be sent to Colin Klein, Philosophy Department, MC 267, 601 S. Morgan Street, University of Illinois at Chicago, Chicago, Illinois 60607–7114. Email: cvklein@uic.edu

should make a "distinctive contribution to the practice of cognitive science" (p. 18).[1] Rupert is appropriately broad about what would count as empirical progress; it could come in many forms, "ranging from innovative experimental designs that produce new effects (or otherwise impressive data) to a more unified or explanatorily powerful theoretical interpretation of existing results" (p. 18).

In chapter 3, Rupert outlines his distinctive systems-based approach, which he argues fits the bill. Rupert offers as a condition of demarcation that "a state is cognitive if and only if it consists in, or is realized by, the activation of one or more mechanisms that are elements of the integrated set members of which contribute causally and distinctively to the production of cognitive phenomena" (p. 42). The notion of an integrated cognitive system is cashed out in section 3.2.2. Simplifying, Rupert offers a two-step process for finding the integrated cognitive set. Start with the set of all mechanisms that have *ever* contributed to a subject's intelligent behavior.[2] This set will include many things, including all of the external resources cited by HEC. First cut: rank sets of mechanisms based on how often they are co-employed in the performance of cognitive tasks.[3] Find a suitable gap, and throw out every set below it. Second cut: count how many times each mechanism appears in the remaining sets. Again, find a gap in the ranking, and throw out all of the mechanisms that fall below the gap. The resulting set of mechanisms will be those that are "highly interdependent and heavily co-employed" (p. 43). Members of this set, and this set alone, should be counted among the cognitive.

I have a few worries about Rupert's particular formulation. Step one would seem to undercount the contribution of inhibitory mechanisms like executive control: while they may contribute to many tasks, they need not be frequently co-active with any *particular* set of cognitive mechanisms.[4] This could be plausibly avoided by tweaking the sort of conditional probabilities employed.[5] Further, as Rupert notes, his condition doesn't actually rule out everything he'd like it too. Ambient light contributes, frequently, to my perception. Rupert suggests that we add a third cut, one which screens off mechanisms that contribute to cognition only by their effects on a single mechanism (like perceptual input systems) [p. 42]. Proponents of HEC should balk at this: whether perceptually mediated mechanisms can be part of cognition is precisely what's at issue in the debate. Such a criterion might be independently motivated (see for example Weiskopf, in press). Lacking independent motivation, Rupert might do

[1]The first two *desiderata* are that a demarcation criterion should amount to more than relabeling of an existing successful practice, and that it should not multiply entities beyond necessity. While these are rhetorically useful, the third *desideratum* is clearly the weight-bearing one — as Rupert notes, a proponent of HEC who showed a good empirical payoff over existing practice would automatically satisfy the first two criterion: the relabeling and expansion implied by HEC would be justified by their contributions to increased empirical success.

[2]Rupert casts the net wider: mechanisms, capacities, abilities, and so on (p. 42). I'll stick with "mechanisms" for simplicity.

[3]More precisely, for each mechanism m, calculate the conditional probability of its employment $p(m|S)$ relative to each subset S of the whole. Then rank order each set mUS by that conditional probability.

[4]See Klein (in press), for an argument that might be a plausible neural architecture.

[5]One possibility: rather than looking for sets with a *high* probability of co-contribution, as Rupert suggests, one could look instead for sets whose co-activation *changes* the probability of use of a mechanism over its probability of use conditional on any particular member. That would allow for both increases and decreases in activation.

better to flesh out his notion of *distinctive* contribution to cognitive functioning; ambient light may contribute to many cognitive tasks, but not in any "nontrivial and distinctive" way (p. 19).

Minor worries aside, Rupert claims, plausibly, that the systems-based view meets his three *desiderata* for a demarcation criterion. It doesn't require excessive relabeling of ordinary practice because it doesn't require relabeling at all: cognitive states end up restricted to the organism-bounded system studied by traditional cognitive science. Only there do we find mechanisms that are frequently co-employed; no external mechanism is co-employed across so many experimental contexts as the mechanisms we find inside the organism (p. 46). Ditto for multiplying entities beyond its payoff: since we've stuck largely to the bounds of traditional cognitive science, there is minimal risk of introducing superfluous entities (p. 44). Finally, the systems-based approach has a clear empirical upshot: by focusing on the entities that show up in the integrated system, cognitive science will be guaranteed the study of a rich field of interesting mechanisms that show up, over and over again, in a variety of circumstances (p. 44). Conversely, anything that doesn't end up in the set is best kicked off as an external resource: its contribution to cognition is too tenuous or infrequent for it to be worth including in the central theoretical constructs of cognitive science.

The systems-based approach has several virtues. Chief among them is that it grounds issues about demarcation squarely in the *methodology* of cognitive science. Why cut the cognitive at the boundaries of the organism rather than further out? Well, if Rupert's right, that's where it will be most fruitful to do science: only within the organism do we find a persisting, integrated set of capacities that are responsible for cognitive behavior across contexts. Because these capacities are integrated, they will be worth studying jointly: they work together to do a variety of things. Because these capacities are relatively invariant across contexts, anything we find out will likely be useful in other contexts as well.

This virtue can be approached from another angle. Both sides of the debate have typically focused on the philosophical and empirical stakes involved in calling something cognitive. But the converse question is equally important: what's at stake in calling something an external resource that is merely used by the cognitive system? It's not to say that it can't show up in cognitive science explanations. Most explanations are scientific hybrids: they borrow from other theories to flesh out their explanations depending on target and context (pp. 20–21). So the non-cognitive often shows up in good cognitive science explanations: it's just not part of the distinctively cognitive contribution to the explanation. Nor is it to say that we're justified in ignoring external resources when we build experiments. If Rupert's vision of HEMC is right, the core cognitive system is especially good at interacting with the environment, so discovering its distinctive properties will require careful attention to the environment and the cognitive system's interaction with it (chapters 9 to 11).

Rather, the distinction is this: calling something external means that we are allowed to treat its properties as *given*, at least as far as cognitive theorizing goes. Discovering facts about external resources belongs to some other discipline (or comes from our general background knowledge). A meteor may have wiped out the dinosaurs, but evolutionary biologists aren't thereby obliged to study the properties of meteors: they can leave that to NASA. Similarly, if iPhones are really external resources, cognitive scientists can defer to the folks in Cupertino about their properties; iPhones don't belong to the set of things that cognitive science has to theorize about.

The systems-based view thus makes clear just what's at stake in debates over HEC. If HEC is right, then we've wrongly taken some external properties for granted, and should expect our theorizing to suffer thereby (I'll return to this possibility at the

end). Conversely, if HEC is wrong, then adopting it will waste everyone's time: it will mean poking around cataloging trivialities rather than doing theoretically progressive science.[6] Cognitive scientists should always keep this danger in mind. As Newell (1973) warned nearly forty years ago, coming up with theories about particular phenomena is the *easy* part of science; figuring out how to fit it all together is the really tricky part. The systems-based view says that this unification is most likely to happen if we focus on the integrated, persisting set of capacities that ground cognition in a variety of circumstances; I'm inclined to agree.

Arguments Against HEC

Systems-based view in hand, Rupert relies on two arguments against HEC. The *argument from demarcation* is straightforward: the systems-based view is the only plausible account of demarcation, and since it appears to locate cognition largely within the organism, we should reject HEC (p. 45). The *argument from empirical success and methodology* claims that the systems-based view can account both for the success of contemporary cognitive science and for the experiments that impress proponents of HEC; HEC-friendly approaches, by contrast, either fail to accomplish the first goal or else end up less simple, conservative, or explanatorily powerful than the systems-based view (p. 46).

While the arguments work in tandem, the first three chapters are largely devoted to defending the argument from demarcation, while chapters 4 through 8 rely heavily on the argument from empirical success. I'll consider both in turn.

I've focused above on the positive argument for the systems-based view, but Rupert also spends considerable time arguing that other proposed principles fail to meet his *desiderata* for a demarcation principle. Much of this is well-trodden ground, though no less impressive for that. Identifying the cognitive as whatever makes a causal contribution to cognition includes far too much — distant stars, Aristotle, and anything else that might impact our perception; so either it involves gratuitous relabeling or else massively expands our ontology without obvious payoff (pp. 19–20). Ditto for labeling as cognitive anything that shows up in cognitive science explanations (p. 20) or which is necessary or sufficient for cognitive functioning (pp. 22–25). Nor will it help to appeal to organism-centered but extended cognition (p. 47), or to growing and shrinking systems (p. 49), or to external information that is consciously endorsed (p. 28): each of these intuitively privileges the persisting integrated core system, and so appear to be nothing more than needless relabeling of the successful practice captured by the systems-based approach. I find little to object to here: Rupert is clearly right that many of these proposals are either straightforwardly inferior to the systems-based approach, or presuppose something like the systems-based approach to motivate their expansion.

The argument from empirical success, on the other hand, is easy to misinterpret. Throughout most of parts two and three, Rupert employs it in a characteristic way: he takes some empirical results that impress HEC, shows that HEMC can also accommodate them (typically by appealing to internal representations of external states) and concludes that, on balance, we should prefer HEMC. Consider his discussion of

[6]Clark suggests that cognitive science has no reason not to move beyond the stable, persisting system and thereby "let a thousand flowers bloom" (2007, p. 170). It's worth remembering that this slogan is a double-edged sword: it has its origin in the Mao-era Hundred Flowers Campaign, which caused considerable chaos with little benefit for most involved (Spence, 1991, chapter 20).

the work of Gray, Sims, Fu, and Schoelles (2006) on the apparent equivalence of internal and external information stores. Clark has cited this as an example of the impartiality of the cognitive system to the location of stored information, an apparently HEC-friendly result (Clark, 2007). Rupert interprets it as showing something different: "that when there is no great cost in terms of time, the cognitive system will use resources beyond its boundary" (p. 103). This use of external resources is grounded in the use of internal cognitive capacities that form part of the core system: memory, perceptuo-motor routines, and so on. Gray et al.'s results don't favor HEC over HEMC; on grounds of conservativism, simplicity, and explanatory power, we should therefore prefer HEMC (p. 105).

Here, the proponent of HEC might object in two ways. First, it's not obvious that HEMC actually wins out on criterion of simplicity: HEC appeals to a single process (retrieving information) that can cross the organism's boundaries, while HEMC must appeal to a series of more complex processes (cf. Clark and Chalmers, 1998, p. 13; Sprevak, 2009, p. 524). Second, this argument strategy seems to unduly privilege HEMC simply because it is the more conservative position. Surely part of the motivation for HEC, though, is the thought that traditional cognitive science is deficient; why give HEMC special weight just because it is more closely aligned with traditional cognitive science?

The first objection is misguided, and Rupert's account shows why. Simplicity, conservativism, and explanatory power are virtues of whole theories, not individual explanations (Fodor, 2009). The systems-based view correctly places the focus on whole theories: it is only, Rupert argues, when we focus on the whole integrated system that such virtues emerge. Individual HEC-friendly explanations thus might be simpler on a case-by-case basis, but lose out overall because they explain piecemeal what the traditional view explains in a unified way.

Consider memory. Rupert notes that HEC could propose a single type of memory that crosses organismic boundaries, functionally differentiated by its ability to access information. This might look simpler than a dual theory that proposes both intrinsic and extrinsic stores. But, as Rupert notes, this coarse-grained version of memory will typically *lose* on considerations of overall explanatory power: it is unlikely to explain phenomena like negative transfer that depend on the fine-grained details of internal capacities (pp. 97–99; Rupert, 2004). Thus the nonextended view will win out on explanatory power and simplicity overall, despite giving a more complex explanation of certain HEC-friendly cases.

This observation goes a long way towards answering the second objection. The first premise of the argument from empirical success does not, note, rely on principles of conservativism. Nor does it require us to think that cognitive science has been especially successful. Rather, it says that cognitive science has had *some* empirical successes, and one job of an alternative HEC-friendly principle of demarcation is to explain why its opponent has been successful. The problem with an extension-friendly functionalization of memory is that it gives up on explaining these successes: the more general functional type lacks the detail necessary to ground interesting generalizations about memory, and it gives us no payoff beyond explaining HEC-friendly cases (which HEMC can explain as well). *This* is the point where the theoretical virtues come in: we should, all things considered, prefer a well-established theory to one that wins on simplicity but abandons a lot of explanatory power.

The argument from explanatory success drives many of the arguments in Part II, though Rupert usually also relies on more specific objections against HEC. I can't do justice to the breadth of Rupert's book and his arguments; he covers essentially every

interesting argument for HEC, and shows either that they can be embraced by HEMC or else that they fail on independent grounds. These include arguments from developmental systems theory and the external scaffolding of language (chapter 6), the phenomenal experience of extension (chapter 8), and dynamic systems theory (chapter 7). For each, Rupert does a terrific job of cataloging and clarifying various arguments in the field, and showing that each is either implausible or else does not conflict with a HEMC-based alternative.

The Positive Alternative

As Rupert notes, much of his argument is explicitly comparative: it's not simply that HEC fails to account for this or that phenomenon, but that it offers little over an organism-centered systems-based approach (p. 241). Part III lays out some of the details of HEMC, the positive proposal that Rupert consistently adverts to. Indeed, if you're approaching the book with a skeptical eye, I'd recommend reading Part III before Part II; its goal is primarily to show that an embedded approach has more explanatory resources than is typically supposed by proponents of HEC, and so *can* account for many of the phenomena covered in Part II.

The embedded approach tries to model cognitive phenomena using fewer, less complex mechanisms than traditional cognitive science uses (p. 180). It is still computationalist in spirit, but allows that computation processes might be time-sensitive (p. 190), involve egocentric and action-oriented representations (p. 196ff), represent environmental affordances (p. 203), and be flexible in the nature and origin of atomic concepts and basic computational architecture (p. 209ff). Proponents of HEC have tended to take the shortcomings of traditional computationalist approaches as shortcomings of a computationalism as a whole. Rupert is right to note that computationalism is a big tent (p. 188), and so flexible enough to accommodate many of the results that impress fans of HEC.

The hypothesis of embedded cognition thus allows for the possibility of substantial revision to traditional cognitive models. This fact is, I think, partly in tension with the argument from empirical success — if taken far enough, the possibility of revision threatens to undercut the traditional results that Rupert wants HEC to explain. That said, HEC does not depart radically from traditional cognitive science. In particular, it allows that traditional sorts of explanation might coexist alongside more embedded ones: some capacities might require relatively abstract, subject-neutral representations, while others might require time-sensitive, action-grounded, egocentric representations.

HEMC is thus an extremely flexible position. That is as it should be, and points to a final virtue of Rupert's systems-based criterion for demarcation. Nearly every party in the debate on HEC agrees that we should look to our best cognitive science to demarcate systems. But what *part* of science should we look to? Most demarcation accounts are *content-based*. They look at the content of particular cognitive theories, extract from them some property that is distinctive of cognition, and then restrict or expand the cognitive to just those things that have this mark. Some of these are HEC-friendly: the mark of the cognitive is its coarse-grained functional profile, or its causal contribution to cognitively distinctive tasks, or simply the fact that it appears in good cognitive explanations. Others are HEC-unfriendly: the mark of the cognitive is its fine-grained functional profile, or its distinctive degree of informational integration (Weiskopf, 2008), or its non-derived content (Adams and Aizawa, 2008) and intensionality (Fodor, 2009; Horgan and Kriegel, 2008).

Rupert's criterion, in contrast, is *method-based*: it says that we should look at the *techniques* of successful cognitive science, and make sure that our method of demarcation preserves the successful use of those techniques. Both content- and method-based demarcation are initially plausible: both look for a mark of good science, and say that we should keep doing that. But I think the flexibility of HEMC shows a reason to prefer a method-based account over any content-based account, HEC-friendly or unfriendly.

Most obviously, method-based demarcation is more ecumenical. Cognitive science is still a young discipline. Our confidence in any proposed mark is low; hence, the plausibility of particular mark-based views will depend a lot on what you think of HEC in the first place. A particular kind of content — traditional computationalism, say — may have been popular for historically contingent reasons (it was mathematically tractable, or it fit better with the scientific *zeitgeist*). Hence, any content-based view runs the risk of sounding unreasonably conservative: fans of HEC can (and do) say that the theories that did work need not be the best guide to what will work in the future. Method-based demarcation, by contrast, is more flexible. Need egocentric representations in your theory? Add them in, regardless of whether they've been useful before. So long as you're using them in the same *way* as previously successful theories used their representations — that is, to track the properties of persisting, integrated systems — you can have some confidence in future empirical success. Hence the flexibility of HEMC: it can easily accommodate new experimental results, so long as they track the empirically useful persisting system.

The point can be put in a slightly different way. Though this is rarely made explicit, HEC defenders really fall into two camps: the *radicals* and the *nihilists*.[7] The former believe that the cognitive does have boundaries, and we've gotten them wrong. The latter claim that there is no useful fact of the matter at all: that "what is thought to be a 'system' or sub-system is, after all, just a convention" (Chemero and Silberstein, 2008, p. 131).

Content-based approaches answer only the radicals. The nihilists, on the other hand, can cheerfully admit that studying things with a content-based mark is useful in some contexts, not useful in others, and so need not be privileged. There is thus a real risk of the debate turning into an argument over mere terminology. Suppose we find that some mechanism plays a critical and integrated role in cognition, so critical that leaving it out would seriously hamper the empirical advance of cognitive science, but it lacks a proposed mark. What follows from that? Why shouldn't the HEC proponent concede, call this new mechanism Cognitive*, propose a new, more fruitful Cognitive* Science, and move on? *Ex hypothesi*, leaving out the mark-free mechanism would hurt Cognitive Science, and so give us a reason to prefer Cognitive* Science.[8] What justifies a mark is either its empirical fruitfulness or something else. If the former, we might as well (as the systems-based approach does) just demarcate by empirical fruitfulness and abandon content-based demarcation; if it's something else, it's tough to see why cognitive scientists ought to care about it.

The systems-based view, by contrast, answers both the radical and the nihilist. To the former, it provides a compelling reason for drawing the line at the organism's boundary. To the latter, it provides a compelling reason to draw a line *at all*: things within the line have been most fruitful to study, and so we should expect them to con-

[7]Rupert (pp. 52ff, 205ff) does not make this distinction explicitly, but responds to what I'm calling nihilist arguments in several places.

[8]Rupert makes a similar point in his review of Adams and Aizawa's book (Rupert, 2010).

tinue to be so in virtue of forming an integrated and persisting system. Giving that up, in contrast, should be done with reluctance. As Rupert notes, most scientific experiments are written up as if there is a single subject who enters a room, does some stuff, perhaps comes back a few weeks later and does something different, and so on (p. 205). Ignoring this simple fact and moving to nihilism would appear to undercut the very procedure of most of our good science.

Final Thoughts on the Future of HEC

Rupert's book provides a powerful and sustained argument against HEC. Nevertheless, the falsity of HEC is a contingent matter, and we could in principle get results that support it. Rupert suggests a few in passing. I think he actually leaves the door open wider than he thinks. I conclude by making a methodological suggestion about how HEC might be supported even given Rupert's criticisms.

First, I'll take as a given that there *are* scientifically interesting examples of what Clark calls "hybrid" systems — that is, systems with parts that cross interesting ontological boundaries. Sociology and economics (Wilson's [2004] "fragile sciences") are obvious examples. Or consider a more homely one — traffic science.[9] Traffic science produces models that include cars, people, stoplights, road banking angles, lane numbers and widths, and so on. These don't have anything obvious in common. That shouldn't matter: part of the advantage of a method-based approach to demarcation is that it can handle these kinds of hybrid systems if they arise. So long as the relevant bits are integrated in a theoretically fruitful way, we can study them.[10]

How could a fan of HEC show that the organism is part of a hybrid cognitive system? Not by appeal to individual experiments. Rupert's strategy shows that any individual experiment can easily be accommodated by the more conservative HEMC. Trying to defend HEC on a case-by-case basis thus looks like a fruitless exercise in cataloging minor results, in the same way that Creationism often looks like a quixotic attempt to undermine a well-established coherent theory by poking around in minor, usually easy-to-accommodate, particulars. Nor will it do to appeal to very general principles, like "Humans are very good at dealing with tools" — no one denies that, and HEMC is designed to accommodate it without extending the cognitive into the world.

If there are HEC-friendly models, then, they will have to lie in the mid-range: general enough to give a theoretically unified account of a wide variety of our interactions with the environment (satisfying Rupert's demarcation criterion) but specific enough to make useful predictions in individual cases. Note that traffic science provides precisely these sorts of mid-range models: it unifies a wide variety of traffic phenomena, while giving enough detail to explain particular jams. Further, no one doubts that humans are in control of cars (at least for a little bit longer), and so that traffic science is, in Clark's term, an "organism-centered" discipline (Clark, 2007, p. 192) It would be a nightmare to drive on a road designed by a psychologist, though. Knowing things about road banking, lane width, stoplight placement, and so on is absolutely crucial

[9]For a good introduction, see Gazis (1974). Some models of traffic are built using the dynamic systems framework favored by many proponents of HEC; see for example Prigogine and Herman (1971).

[10]It might be objected that "traffic" does not pick out unified systems, and so the analogy is not apt. I doubt this. First, traffic is not just (for example) a bunch of cars near each other: the same distribution of cars in a stadium parking lot does not count as traffic *during* the game, though it might *afterward* — the dynamic context is clearly important (thanks to Lara Buchak for the example).

for studying traffic. We *could* try to leave these out and do traffic science by appeal solely to the representational states of drivers. That would result in needlessly complex models at best, accidents at worst: the appeal of a hybrid science is that, by crossing ordinary boundaries, we get cleaner, more fruitful models.

Rupert suggests that these sorts of models might be a possibility for cognitive science, though he finds no evidence of them in practice. In his chapter on dynamic systems theory, for example, he notes that dynamic models that treat the order parameter of an extended system as a control parameter of a part (and vice-versa), and that do not neatly decompose into simpler systems, would fit the bill (p. 149). He argues plausibly that all existing dynamic models fall short of this criterion (p. 137ff). Nevertheless, if there *were* such a system, it would be clear why the organism-bounded system would be an inappropriate place to cut. To return to a theme from the first section, in such tightly coupled cases we couldn't treat the external as a *mere* resource the properties of which are taken as a given: the interaction of the organism with the external changes both in such a tight way that good science is obliged to treat them as a package.

Are there any good examples of such mid-range systems? Here, I admit a soft spot for Gibsonian defenses of HEC.[11] The Gibsonian picture, broadly construed, is one in which there are certain persisting features of the environment (the affordances) with which we regularly and reliably interact by picking up on perceptual invariants in the environment (Reed, 1996). The interaction between the organism and affordances is typically supposed to be a tight one. Further, ignoring the structure of affordances would be bad in the same way that leaving stoplights out of traffic models would be bad — it would lead you, all things being equal, to build bad models (Chemero and Heyser, 2005).

I have the same qualms about Gibsonian models that most people do.[12] But I think the form of such accounts is worth taking seriously. They are mid-range models precisely because they contain *abstract* properties (affordances, perceptual invariants) that are environmental but still integrated tightly into cognitive models.

Rupert gives several arguments against extension via these abstract properties (p. 55ff); I'm not sure any of them really have bite. He notes that the arguments for HEC typically presuppose tight causal coupling with finely individuated external properties (p. 55). That's not a problem: with the right metaphysics, a fan of abstract properties can say that the causal interactions are always with fine-grained properties, while scientific *models* gain generality by incorporating abstract properties that cover the fine-grained cases.[13] Rupert

---

[11]See for example (Chemero, 2009), which combines a Gibsonian and a dynamic-systems defense of HEC.

[12]Briefly: the part about perceptual invariants seem to be compatible with traditional vision science, and the use of such information to pick up affordances would seem to either require a rich internal structure like that postulated by HEMC (in which Rupert's arguments would seem sufficient to block it) or else an implausible version of associationism. Rupert shares these concerns (pp. 196ff, 203).

[13]Compare: traffic models typically include "number of lanes" as a variable, even though individual cars are always driving on a particular road with a specific number of lanes. That's unremarkable though: science gets its power precisely through this sort of abstraction (Klein, 2009). So, for example, the abstract property could be a determinable, particular determinates of which enter into tight causal interactions. Or the property could be a second-order property, particular first-order instantiations of which are the more specific causes of behavior. Or we could eschew talk of properties and just call them abstract *descriptions* that are obligatory components of the simplest mathematical models of behavior.

objects that external abstract properties are not representational (p. 57). That might be a problem if we were working with a content-based version of demarcation, or if the abstract properties were meant to *replace* internal mechanisms rather than merely augment them. Rupert isn't committed to the former, though, and the view under consideration does not assume the latter.

Finally, Rupert notes that there must be a mechanism by which these abstract properties affect the organism, and that's plausibly good-old fashioned representation (p. 56). But remember, the question is not whether an internalist account *could* accommodate these things, but whether the resulting model would be simpler and more powerful than one which simply treated external abstract properties as explanatory primitives. Again, compare traffic: we *could* tell a story that involved representations of lane widths and stopping distances and so on, but what would it get us? Rupert suggests that external abstract properties ultimately will be less integrated than the core internal properties (p. 57), and so HEMC will (again) win out on grounds of explanatory power and simplicity. That may well be true in many cases. But it's not obvious that this is something Rupert can take for granted. The existence of things like traffic science shows us that, at *some* level of abstraction, there will be useful and integrated scientific systems that include individuals but extend beyond them. So if there are abstract external properties that are theoretically important, persistent, and well-integrated into our normal cognitive mechanisms, then the details really will matter — the argument from explanatory success does not have the same easy application that it does when Rupert discusses individual experimental results.

Of course, that's all very general; I think Rupert is probably right that no existing HEC proposal fits the bill. Nevertheless, it seems to me that mid-range models involving abstract external properties are precisely where HEC should be looking. That said, all of this is a good (if perhaps unintended) feature of Rupert's book: it makes clear just what those who disagree with him would have to do to overcome his objections.

Rupert's work is impressive both in its scope and its depth. To conclude, I'll give it the highest praise that I can think to give a book in philosophy: before I read it, I was on the fence about HEC. The book convinced me. It is excellent work, and one that should be read by anyone interested in the debate.

References

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

Chemero, A. (2009). *Radical embodied cognitive science.* Cambridge, Massachusetts: MIT Press.

Chemero, A., and Heyser, C. (2005). Object exploration and the problem with reductionism. *Synthese, 147,* 403–423.

Chemero, A., and Silberstein, M. (2008). Defending extended cognition. In B. Love, K. McRae, and V. Sloutsky (Eds.), *Proceedings of the 30th Annual Meeting of the Cognitive Science Society* (pp. 129–134). Austin, Texas: Cognitive Science Society.

Clark, A. (2007). Curing cognitive hiccups: A defense of the extended mind. *The Journal of Philosophy, 104,* 163–192.

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

Fodor, J. (2009). Where is my mind? *London Review of Books, 12,* 13–15.

Gazis, D. (1974). *Traffic science.* New York: John Wiley & Sons.

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.

Horgan, T., and Kriegel, U. (2008). Phenomenal intentionality meets the extended mind. *The Monist, 91*, 347–373.
Klein, C. (2009). Reduction without reductionism: A defence of Nagel on connectability. *The Philosophical Quarterly, 59*, 39–53.
Klein, C. (in press). Redeployed functions versus spreading activation: A potential confound [Commentary on M. Anderson]. *Behavioral and Brain Sciences*.
Newell, A. (1973). You can't play 20 questions with nature and win. In W. Chase (Ed.), *Visual information processing* (pp. 283–308). New York: Academic Press.
Prigogine, I., and Herman, R. (1971). *Kinetic theory of vehicular traffic*. New York: Elsevier.
Reed, E. (1996). *Encountering the world: Toward an ecological psychology*. New York: Oxford University Press.
Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *The Journal of Philosophy, 101*, 389–428.
Rupert, R. (2009). *Cognitive systems and the extended mind*. New York: Oxford University Press.
Rupert, R. (2010). Systems, functions, and intrinsic natures: On Adams and Aizawa's *The Bounds of Cognition*. *Philosophical Psychology, 23*, 113–123.
Spence, J. (1991). *The search for modern China*. New York: W.W. Norton & Company.
Sprevak, M. (2009). Extended cognition and functionalism. *The Journal of Philosophy, 106*, 503–527.
Weiskopf, D. (2008). Patrolling the mind's boundaries. *Erkenntnis, 68*, 265–276.
Weiskopf, D. (in press). The Goldilocks problem and extended cognition. *Cognitive Systems Research*.
Wilson, R. (2004). *Boundaries of the mind: The individual in the fragile sciences*. New York: Cambridge University Press.