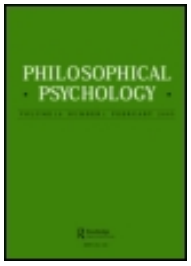

Philosophical Psychology

Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/cphp20

Ten questions concerning extended cognition

Robert A. Wilson

Published online: 13 Aug 2013.

To cite this article: Robert A. Wilson , Philosophical Psychology (2013): Ten questions concerning extended cognition, Philosophical Psychology, DOI: 10.1080/09515089.2013.828568

To link to this article: http://dx.doi.org/10.1080/09515089.2013.828568

Philosophical Psychology, 2013
http://dx.doi.org/10.1080/09515089.2013.828568

Routledge
Taylor & Francis Group

Ten questions concerning extended cognition

Robert A. Wilson

This paper considers ten questions that those puzzled by or skeptical of extended cognition have posed. Discussion of these questions ranges across substantive, methodological, and dialectical issues in the ongoing debate over extended cognition, such as whether the issue between proponents and opponents of extended cognition is merely semantic or a matter of convention; whether extended cognition should be treated in the same way as extended biology; and whether conscious mental states pose a special problem for the extended mind thesis. Lethocerus, the giant water bug, will be a recurrent reference point for much of the discussion.

Keywords: Consciousness; Extended Biology; Extended Cognition

Introduction

The extended mind thesis says that the mind and the cognitive processes or states that constitute it physically extend beyond the boundary of the individual agent. The thesis had precursor forms in Harman's (1988) *wide functionalism* and in my *wide computationalism* (Wilson, 1994b), and was so-named by Clark and Chalmers (1998) in what has become one of the most cited papers in philosophy in the past 25 years. Proponents of the extended mind thesis sometimes talk of extended *cognitive systems*, and these are extended in just the sense specified above: the cognitive processes or states constituting such a system physically extend beyond the boundary of the individual agent. The extended mind thesis and extended cognition are radical forms of externalism about the mind.

As forms of *externalism*, they contrast with *internalist* or *individualistic* views of the mind and cognition, which hold (in one influential formulation, Fodor, 1987, chapter 2) that cognition supervenes on the intrinsic, physical properties of the individual. Individualists allow that cognition can be causally influenced, perhaps in deep and

Robert A. Wilson is professor of philosophy and of educational policy studies at the University of Alberta. Correspondence to: Robert A. Wilson, 2-40 Assiniboia Hall, University of Alberta, Edmonton, Alberta, Canada T6J2E5. Email: rwilson.robert@gmail.com

important ways, as recent proponents of the *hypothesis of embedded cognition* emphasize (Rupert, 2004, 2009), but that cognition itself takes place solely within the head.

As *radical* forms of externalism, these views also contrast with earlier, traditional forms of externalism, which challenged individualism by arguing (e.g., Burge, 1979, 1986) that mental states individuated by representational content did not supervene on an individual's intrinsic, physical properties. Radical externalism does not crucially appeal to representational content, embracing the view, as traditional externalists did not, that cognitive systems themselves physically extend beyond the boundary of the individual.

Rather than rehearsing or recounting the various arguments for or against extended cognition, here I take up some of the broader issues such arguments raise by considering ten questions that I have encountered in extended discussion. The answers provided here derive from my semi-articulate murmurings at the time, massaged into whole sentences with my usual lightning speed, several months later, and then fine-tuned by the passage of time—and by the indefinite delay to the start of the current NHL season.

1. I am a card-carrying externalist about content. But are there good reasons, given that commitment, for adopting a more radical form of externalism, like the extended mind thesis?

Since traditional externalism has gained much support over individualism in the philosophy of mind over the years, arguments levering antecedent commitments to traditional externalism—externalism about mental content—to argue for more radical forms of externalism have a certain kind of dialectical attraction. And there are at least some appealing, general considerations that may be invoked in articulating such arguments. For example, if you're a realist, and you think that, plus or minus a bit, our taxonomic or individuative practices get things right, then externalist explanatory practices should be underwritten by a corresponding externalist metaphysics. Or perhaps one views traditional externalism as an unsteady, halfway-house position, one stabilized as a form of externalism only by accepting the extended mind thesis.

Unfortunately, I don't think there are good arguments from traditional to radical externalism, despite the dialectical promise of such an argument and these kinds of general consideration. In fact, the best arguments for extended cognition don't rely on traditional externalism. They are of two types.

The first of these challenge individualists to defend differential treatment of either internal and external cognitive resources, on the one hand, or cognitive and other systems (such as biological systems), on the other. Such arguments thus rest on one or more considerations of *parity* (Clark & Chalmers, 1998; Wilson, 2004, 2005a). The second identify work in the cognitive sciences that prima facie exemplifies extended cognition (Clark, 2008; Menary, 2010). Both types of argument are naturalistic and also provide a healthy mixture of considerations from the philosophy of science and mind. As such, they sometimes involve offering externalist accounts of key notions, such as functionalism (Wheeler, 2010), realization (Wilson, 2001; 2004, chapters 5–6; Wilson & Craver, 2007), intentionality (Wilson, 2010b), or embodiment (Clark, 2008; Wilson & Foglia, 2011).

2. Isn't the extended mind thesis just a view about what is possible, i.e., minds could extend beyond the body?

No. The extended mind thesis is not simply a view about the possibility of extended minds, and even arguments for that possibility have been formulated with an eye to connecting extended cognition to ongoing explanatory practice in the cognitive sciences. It is true that early arguments for the thesis, most notably the Otto-Inga argument (Clark & Chalmers, 1998), appealed to thought experiments and primarily aimed to loosen the attachment that philosophers of mind and cognitive scientists had to *locationally individualistic* (Wilson, 2004, chapter 4) views of the mind. Yet while looking to establish the possibility of extended cognition, such arguments also reflected a broader interest in offering an alternative, general perspective on the mind and cognition, especially as studied within the cognitive sciences. In a similar fashion, my "Wide computationalism" (1994b) argued that there was nothing about computational individuation requiring that the computational systems driving cognition be individualistic, i.e., that it was possible that there were wide or extended computational systems for cognition. But the paper's argument also included a detailed discussion of two examples of work on cognition—on animal navigation, and on spatial channels theory in form perception—to the effect that the computational systems in each were in fact extended. Likewise, much of Clark's work, from *Being there* (1997) to *Supersizing the mind* (2008), has drawn on a variety of rich and often intriguing examples of research from the cognitive sciences that both support and illustrate "*3E*"—embodied, embedded, extended—views of cognition. Conversely, the connection between explanatory practice and extended cognition has appealed to those working on related topics with a strong naturalistic bent to them, such as moral cognition (Sneddon, 2011), neuroethics (Levy, 2007), and group-level cognition (Theiner, 2008; Theiner & Wilson, 2013; compare Wilson, 2005b).

3. Are you implying that there are at least some cognitive behaviors that simply can't be explained or understood unless the extended mind thesis is true?

No: one should be skeptical about an appeal to "decisive cases," and in general there is a lot of room for both internalists and externalists to move with respect to *any* particular cognitive phenomenon. That said, radical externalists think that for cognition proper to occur, *what's inside* is often supplemented by *what's outside* in a way that the extended mind thesis makes better sense of than does individualism. Some cognitive behavior, such as (to use an example I will return to) problem-solving using external forms of representation, pose a prima facie challenge to individualism and correspondingly seem more readily explained by an appeal to extended cognitive systems.

To clarify the methodological point here, consider the one-line summary of the history of behaviorism: attempts to account for "verbal behavior" and other seemingly mentally driven behavior came to be seen over time to be increasingly desperate and ad hoc (Chomsky, 1959). In the case of behaviorism, one might think

that there was a fundamental problem of the circular and/or empty nature of behavioral analyses (Wilson, 1999). But behaviorists did not always see it that way, and their paradigm for psychological explanation continued on in diminished form. Even with behaviorism's demise as an overarching paradigm for psychology, there remained cases in which a behaviorist perspective cut through excessive mental elaboration unnecessary to account for the phenomenon at hand. The relative success of behavior modification therapy for certain kinds of addiction would be one example.

The point here is not that individualism is in as bad a shape as behaviorism was—it is not and has provided much insight into cognition. Rather, it is that even in a case with failures as extreme as those of behaviorism, there are glimmers of light. To insist that a behaviorist approach could not account for *any* seemingly cognitive phenomena would be mistaken, despite behaviorism's failures as an overarching, reductive research framework. Likewise, even though individualism fails as a similarly constraining research framework, it would be overstatement to claim that it does not possess the resources to account for some specific phenomena: it does. But like the case of behaviorism, there are phenomena for which individualistic explanations become strained, appeal to general principles that can't be defended, and shut off avenues for promising research—in this case, work appealing to extended cognition.

4. Won't traditional views of cognition win out over extended cognition on criteria employed for deciding between theories, such as parsimony, entrenchedness, and generalizability?

No. Those criteria can cut both ways, especially once one takes into account the full range of relevant phenomena, despite the claims made here about both content externalism and the extended mind (Patterson, 1991; Rupert, 2004, 2009; Segal, 1989). For example, although some have thought that individualistic views have greater *causal depth* than content externalist views, content externalist explanations often have both this explanatory virtue and others, such as being more *theoretically appropriate* than their individualistic rivals (Wilson, 1994a; 1995, chapter 8). More recently, I have critiqued Rupert's (2009, pp. 59–60) argument appealing to the relative *temporal grain* of in-the-head and extended cognitive systems (Wilson, 2010c). This argument fails in much the way that did past appeals to putative methodological or theoretic explanatory virtues of individualistic explanations (Wilson, 1995, chapter 4).

Relevant here is my long-standing skepticism about arguments appealing to general features of science, cognition, computation, or representation to argue either for individualism or for externalism as a general constraint on psychological explanation or the boundaries of the mind (Wilson, 1995, 2004). In this respect, I embrace a kind of pluralism. Because I view both traditional externalism and the extended mind thesis as comporting with this pluralistic commitment while individualism does not, the burden of proof on individualists continues to seem to me distinctive.

5. Could an extended cognitive system actually reduce the functionality of a cognitive agent, rather than enhance it?

Yes, though an implicit bias in discussing extended cognition may lead one to think of extended cognitive systems as always functionally enhancing, or always a good thing for the cognizer. They may not be.

One way for cognitive extension to lead to a reduction in functional capacity is through *cognitive clutter*: by adding more bells and whistles to an existing cognitive system, we might well cause it to operate less effectively, or even lose certain kinds of functionality; so-called "smart technologies," for all the benefits they bring, are often used in ways that have this effect for particular tasks (e.g., in driving). Global functional reductions can also result from cognitive extension; sometimes, more is less. Adding a sixth sense (e.g., developed echolocation) to existing human senses that was extremely oxygen demanding, such that its operation interfered with other cognitive functions, could have such an effect. If that sense were not integrated in the right ways with action and other cognitive processes, it might add very little, if anything, to our cognitive functioning but end up imposing significant losses of functionality on other cognitive systems.

Extended cognition has proven useful in thinking about cognitive dysfunction compensated through *reliance on others*, such as often happens in cases of Alzheimer's-induced memory loss and reliance on a trusted other to pick up some of the ensuing cognitive slack (Levy, 2007; Wilson & Lenart, forthcoming). Consider instead a case where the person one relies on is not benevolent but malevolent, or just not a very reliable storage or interactive resource. Here the resultant extended cognitive system could have diminished cognitive capacities relative to the individualistic system it augments. In such a case, extended cognition would not be functionally enhancing, but dysfunctional.

6. Surely the only persisting, unified cognitive system is physically constituted by what is inside the head, or (maybe) inside the body. Aren't putative non-neural cognitive resources simply too ephemeral, derivative, or in some other way secondary to physically constitute such systems?

No, that view of "the non-neural" is mistake. Here, several sorts of parity considerations are relevant. One concerns a comparison of the biological and the cognitive sciences.

Consider *Lethocerus*, the giant water bug. This predatory insect injects its prey with proteolytic enzymes, and then sucks up the liquefied product of this enzymatic action through its proboscis. Here we have a case of *extended digestion*, a process that is not bounded by the physical body of the organism (compare Adams & Aizawa, 2001). Such extended digestion is common in predatory insects and spiders, integrating resources that physically constitute a physiological system crossing the bodily boundary of the organism. Nothing about the resources recruited by *Lethocerus* from beyond its body make them ephemeral, derivative, or secondary to the resources it

uses inside its body in forming an *extended physiological system* (sensu Turner, 2000). While recognizing that much physiology takes place entirely inside the body, biologists and philosophers of biology seem relatively unperturbed by the idea of extended digestion.

In turning from digestion in the biological sciences to (say) *problem solving* in the cognitive sciences, we do not encounter the same ecumenicism about the physically constituent resources or about the resultant, integrative system(s) serving this function. Here, the individualistic tradition in the study of cognition has an unfortunate legacy. By assuming that only the in-the-head resources count as part of the relevant cognitive system, this research tradition, as interesting as it may be for other reasons, tells us little about how people *actually* solve problems.

It is not simply that particular models of problem solving, like the General Problem Solver (Newell, Shaw, & Simon, 1960), are mistaken in detail. It is that the entire approach to problem solving, where problems are well-defined and involve at their heart search through a problem space in some effective manner, fails to connect with real-world problem-solving. Such approaches fail, basically, because the way in which they define the process of problem solving and the strategies they propose, are neither the ones that people use, nor do they reasonably approximate those strategies (Kirsh, 2009; Sawyer & Greeno, 2009).

7. Sometimes I wake up at night and worry that this is all merely a terminological or semantic matter. Is there really a fact of the matter as to whether, for example, one has an extended cognitive system or merely lots of interesting causal interaction between extra-cranial stuff and in-the-head cognition?

A problematic form of underdetermination sufficient to reasonably keep one up at night requires that *every* case could be plausibly explained by individualists and externalists (e.g., as embedded versus extended cognition sensu Rupert, 2009). Fortunately, this is not the case. So, yes, there is an empirical fact of the matter in at least *some* cases, and so no, the individualist-externalist debate and the determinate forms it takes are not merely terminological or semantic. I take there to be a fact of the matter about the kind of cognitive system operating in at least some cases, where a cognitive system is an *integratively coupled system containing cognitive resources* (Wilson, 2010a). When some of those resources are external, the cognitive system is extended.

Return to *Lethocerus* to fix on a non-cognitive, integratively coupled system that is extended, here one containing *digestive resources*. In the *Lethocerus* case, there is causation between the extended part of this system—the injected prey that becomes a container for digestion—and the internal parts of this system—all of the internal digestive organs in *Lethocerus's* body. But here we have *more than* mere causation that amounts to integrative coupling. The injection of proteolytic enzymes into the prey adds functionality to the internal part of the overall digestive system. Without this, *Lethocerus* could not, or would struggle mightily to, digest its prey. And the injected prey comes to be *integrated with* the rest of the digestive system.

If digestion is the physiological process of extracting nutrients for organismic sustenance through the use of bodily produced or controlled agents (such as enzymes), then *that* process itself extends beyond the body of *Lethocerus*. Specifying the biological function—digestion—and the facts about how that function is realized in this particular case, determines that the system performing that function is not wholly contained in the body of *Lethocerus*. Thinking of the beyond-the-body processes simply as external causes coupled to an internal physiological system, and so viewing the digestive system itself as ending at the bodily boundary of *Lethocerus*, as a proponent of a *hypothesis of embedded digestion* might do, fails to correctly identify the system that performs the function of digestion in this type of case.

One could insist that external resources are mere causes to, while internal resources are real parts of, the digestive system, just as one could insist on a *hypothesis of stomach-bounded digestion*, according to which only the stomach is really a part of the digestive system, with the activity of all other bodily parts, such as the intestines and liver, either providing inputs to or accepting outputs from the digestive system. But in both cases, that insistence would seem to fly in the face of the integrative nature of the causal coupling in play.

Consider now a particular cognitive process or function: let it again be problem solving, and let's focus on solving a particular board configuration in the game Rush Hour. The question is whether the system that allows an agent to solve that problem stops at the bodily boundary of that agent, or instead extends into the world beyond that body. On the individualistic view, including the hypothesis of embedded cognition (Rupert, 2004, 2009), the problem is solved by the agent forming internal representations, including those that come from vision, and manipulating and transforming those, with the effects of those transformations resulting in behavior. But the cognitive system that underwrites the problem solving itself is located entirely in the head. On the extended mind view, the problem is solved *interactively* (Kirsh, 2009) in part through repeated visual inspection of the board, actual movements of the cars, and lots of internal processing. These perceptual and motor interactions partially physically constitute the relevant cognitive system.

It is possible that human problem solving uses just an in-the-head cognitive system, with vision providing an initial input specifying the problem to be solved, and bodily action executing the solution to the problem. As a matter of fact, *we don't solve the problem like that*—like a computer might (Kirsh, 2009). Since we can literally *see* certain pathways that are more promising as a way through the problem space, we use the new configurations of the board as resources to solve the problem, especially once the problem becomes even moderately complicated (compare to pen-and-paper multiplication). If the problem requires (say) twenty moves to solve, we proceed by taking new visual fixes on these individually or in clusters, continuing the process of "look, think, move" as part of a problem-solving cycle. We literally use our interactions with our immediate environment to build an integratively coupled system containing external resources that allows us to solve the problem at hand.

8. Doesn't extended cognition imply that we also have "extended agents" or cognizers, since now a part of me is located somewhere beyond my body?

No: extended cognition does not imply extended agency. And that's just as well, since extended agency of this kind would pose several problems that proponents of extended cognition need not burden themselves with.

I have previously defended what I called a *narrow subjects, extended systems* view as a kind of plausible middle-ground position here (Wilson, 2004, pp. 141–143). To see its plausibility as a general view of both biological and cognitive processes, return again to the case of extended digestion. Here *Lethocerus*, bounded by its readily identifiable body, is the agent and the *bearer* or *subject* of this particular process. In this case, the spatiotemporal location of the insect's body provides us with a rough-and-ready idea of *where* the agent is. The philosophical gloss on what makes this rough-and-ready idea kosher is given, in my view, in terms of the *locus of control*: it is the beetle that we can pick out in a picture ("that's the beetle, right there") that is also the locus of control for actions that it initiates, and for physiological processes that contribute to its life. Substantial philosophical reason or empirical evidence (or the former supported by the latter) are needed to engage in the revisionary metaphysics of *things* that would have the beetle itself distributed across distinct spatiotemporal regions. Extending the process of digestion does little to provide such reasons or evidence.

One way to support this as a general position is to consider the case of actions. Does action stop at the skin? The intricacies of action theory to one side, there are good reasons to think not, including the idea that a skin-bound view of action would leave us with an impoverished view of what agents do, one that confuses action with mere movement. Those holding the commonsense view that agents do things that *extend into the world*—like making a cup of coffee, driving a car, or writing a letter—seem to feel no compunction in appealing to a regular bodily-bound agent for such *extended actions*. Here a *narrow agents, extended actions* view is a natural default; the reasons one has for accepting extended actions do not challenge the status quo about the boundaries of agents themselves.

Resistance to the narrow subjects, extended systems view sometimes rests on a crypto-mereological appeal to "parts." To point to a confusion in such appeals, consider the following argument for extended agency that makes an explicit reference to parts:

1. The injected prey is part of the digestive system of *Lethocerus*.
2. The digestive system is part of *Lethocerus*. Therefore,
3. The injected prey is part of *Lethocerus*. Thus,
4. *Lethocerus* is an extended agent.

While there is a reading on which each of 1 and 2 are true, the corresponding argument is invalid because that reading—I think the only one on which those premises are both true—requires different senses of 'part of'. In premise 1, 'part of' means "is physically a part of," and so we can talk about the physical stuff that realizes the digestive system, a physical whole, as extending beyond the body in terms

of the notion of wide realization. But that mereological understanding of 'part of' makes little sense in premise 2. In premise 2, 'part of' means something more like "belongs to"—at least that is all that someone inclined to distinguish between systems and agents should accept. Precisely the same is true of an argument stated in terms of the relationships between external cognitive resources (such as a Rush Hour board and pieces), an agent's problem solving system, and that cognitive agent him/herself.

If this is correct, then such arguments equivocate on 'part of'. For this reason, the implicit reliance on the transitivity of part-whole relations in the argument is problematic. This raises further complexities for proponents of constitution views in metaphysics, and for related discussions of realization and mechanism that appeal to the transitivity of 'part of' (Wilson, 2007, 2009).

9. Even if we grant that perceptual or cognitive systems can be extended, shouldn't we reject the idea that this tells us anything about the corresponding experience?

No: once you make that concession, individualism about perceptual or cognitive experience faces an uphill battle. In discussing a draft of my "Extended vision" (2010a), Ned Block suggested otherwise, saying:

> Suppose that you are right that conscious visual mechanisms and representations sometimes are parts of a whole system involving things outside the brain, and that so doing extends the existing functions of the components …. Suppose you are right that some structures outside the head have a claim to be classified as part of the visual system. None of these suppositions is even slightly relevant to the question of whether the minimal physical basis of conscious vision extends outside the head. I am sure that everyone would agree that conscious visual representations and mechanisms are linked via integrative coupling to unconscious visual representations and mechanisms inside the brain, and no one takes that to show that unconscious visual representations and mechanisms inside the brain are part of the minimal neural basis of conscious vision. You have added an outside-the-head set of unconscious visual mechanisms, but that doesn't change the logic at all. (personal communication, February 2009)

The idea here is that extended vision per se has *no* implications for what we say about visual experience; Block takes the existence of extended visual systems to be compatible with his own view that the minimum supervenience base for visual experience is the brain, and nothing larger.

I think that this is mistaken. In my initial response, I said the following:

> Well, I'd say that they are at least slightly relevant, for a few reasons: (a) if one allows that vision in general can be extended, then it seems that the onus is on those who want to argue that there's something special about conscious vision that makes it narrow to explain just why that is; compare those who allow that in general realizers of mental states can be made of metal rather than meat but who then say, "ah, but not the conscious mental states"; (b) visual extension carries with it extended (or wide) realization of visual processing, and while that itself doesn't entail extended realization of consciousness (or, for that matter intentionality, or any particular

> mental property or feature), it does suggest a metaphysics more resistant to craniocentric answers to questions about the realization base for any particular mental property. ...

This response can be strengthened by defending a substantial view of the nature of consciousness as a mental property, and of the realization of conscious experience. What follows is a whiff of such a defense.

First, there are different kinds of properties of cognitive systems (compare biological systems). In particular, consider the following trichotomy of properties that I will illustrate with properties of visual systems:

- Properties of only whole cognitive systems (examples: enables recognition by a perceiver; realizes prey detection; produces rapid aversive behavior)
- Properties of only parts of cognitive systems (examples: transduces light; computes distance from retinal disparity; responds to specific wavelengths of light)
- Properties of both whole cognitive systems and their parts (examples: operates with a specific speed; is damaged; is realized by particular physical stuff; processes visual information)

If this is correct, then the question is where consciousness fits in this trichotomy. My view is that it belongs in the first of these categories. If that is correct, then Block's claim—that showing that visual systems are extended is *not relevant* to showing that visual experience is extended—is false. In fact, it would be *highly relevant*, since it would license the inference from "X is extended" to "the conscious aspects to X are extended," where X is some particular extended (visual) system.

Why think that consciousness is a property only of whole cognitive systems, in this case, of whole visual systems? Consider, first, an *internal* system that results in experience, the nociceptive system and *the experience of pain*. Pain is not realized just in any part of the nociceptive system—the C-fibers, the A-fibers, the nociceptors, limbic system, somatosensory cortex, etc., but in the system as a whole. If you just had those parts in whatever state you like, they would not instantiate the experience of pain. Not unless you think there could be the experience of pain *without a perceiver of pain*, since it is possible to have *any* part of the system in whatever state you like without there being a subject or bearer of pain. In this respect, pain contrasts with properties like *spiking rate*, which is primarily a property of only some parts of the nociceptive system (though one that could be calculated derivatively for the whole system), and *being myelinated*, which is a property only of the A-delta fibers in the nociceptive system.

When people talk of pain being *located in the somatosensory cortex* or being *experienced in the brain*, they are at best talking of pain's *core* realization, not its *total* realization (sensu Shoemaker, 1981). Importantly, only the total realization suffices for the property instantiated; a core realization does not. The misimpression that just by stimulating a localized part of the brain we can induce pain—or indeed any folk psychological state—is one product of this confusion.

Consider now the sort of properties of visual systems that I have listed as being *whole system only* properties: enables recognition by a perceiver; realizes prey

detection; produces rapid aversive behavior. These are properties, respectively, that the recognition, prey detection, and "where/how" system might possess in a given case. They are not, however, properties also of any particular parts of those systems, at least not any of the readily identifiable parts, such as those picked out by physiology, anatomy, or by psychological means, such as via the notion of modularity. This includes the brain as a part, although, strictly speaking, the brain as a whole is not a part of any of these systems; it is only specific parts of the brain that are such parts. In each case, the functioning of these parts might *contribute to* the enablement, realization, or production of the respective property. But these are not properties of those parts any more than folk psychological states are properties of parts of agents, rather than of those agents themselves.

My claim is that consciousness—here visual consciousness—belongs with those properties. One way to argue for this claim would be to develop a general account of the trichotomy between whole system only, part only, and both whole system and part properties, or at least an account general enough to cover either visual, perceptual, or cognitive systems (in increasing order of generality). Then argue that visual consciousness shares in the property(ies) that mark off the whole system only parts. Lacking such a general account, I rely on a technique I learned at grad school: to sneak up on such an account.

The kinds of properties that are whole system properties tend to cluster around the *function(s)* of the respective system. If those functions are performed only by the whole system, and not by any particular parts of that system, then we can understand why properties that cluster around or are closely related to those functions are also whole system properties. These are also properties that the *subject* or *bearer* of those properties has, rather than simply being possessed by the parts of the cognitive systems those subjects or bearers also possess. Although cognitive systems and such subjects or bearers are not identical, intuitively they are more like one another than either is to parts of a cognitive system.

To recount how I see the dialectic here: those, like Block, who think that you can grant everything that a proponent of extended vision argues for *except the claims about the extension of visual experience* are in a difficult position. My argument for extended vision is premised on a claim about the global function of visual systems, with subsequent claims that visual systems are embodied, and because of the *way* in which they are embodied, are also extended. Someone in Block's position concedes all of this, but then wants to draw a line in the sand over extended visual *experience*. Were visual experience a kind of icing on the cake of visual processing, icing that some cakes (like ours) have, but that others (like those of insects) do not, resistance here might have a dialectic anchor. *But that's precisely part of what is given up*, I take it, in adopting the position that Block wants to maintain.

Minimally, there is a burden of proof facing individualists to provide a reason for thinking that visual experience or consciousness is a property of parts of the visual system, rather than the whole visual system or the perceiver. Block himself seems to think that this is obvious—consider his comment above on nonconscious, internal processing. I believe this view is not only not obvious; it is likely indefensible.

10. Parity between biological and cognitive systems only takes you so far. After all, isn't cognition tied up with the nature of agency, individuals, and subjecthood more intimately than is biology, making the extended mind thesis harder to defend?

Probably not, in both cases.

First, as a way to make both question and answer less telegraphic, consider the following thought experiment, a version of which was first suggested to me by Christian Lee. *Lethocerus* injects its prey, walks away, and then is squashed. On my view, at that point, there is (partial) digestion going on—in the injected prey—but *no agent of digestion*. I think that this is a possible, indeed plausible, way to describe what happens—a case that I might call *digestion interrupted*.

Contrast this with a putative case of *cognition interrupted*: Otto remembers by writing some stuff down in his notebook, but then suffers a series of unfortunate events that leave us, alas, with *no Otto*. If we treat the digestion and cognition cases parity-wise, then this should be a case where *there is cognition going on—more specifically, remembering—but no cognizer—more specifically, no rememberer*. For as in the squashed bug case, the external part of the process has begun, even though the agent or subject of that process meets an untimely end. But cognition without a cognizer—here remembering without a rememberer—is deeply counterintuitive in a way that digestion without an agent of digestion is not. This suggests an important disanalogy between the biological and cognitive cases—and a parallel limit to parity appeals—as well as indicates how agency and cognition are more intimately related than are agency and (say) biology.

While I agree that there is an asymmetry here, I don't think that it concerns cognition per se. To see this, consider a case that involves some form of *implicit cognition*, or the unconscious cognitive processing that is commonplace in computational accounts of cognition. Think, in the Otto case, of the notepad as being unconsciously accessed and used as a memory storehouse device. Here the *cognition interrupted* option does not seem to me counterintuitive at all; even those for whom it remains somewhat counterintuitive typically acknowledge that the shift from conscious to implicit cognition weakens the degree of counterintuitiveness. If that is correct, then this suggests that the divide here, if there is one, is between *conscious phenomena* and everything else, not between cognition per se and (say) biology.

But I would also ask the following parity-motivated question: can we have cases of cognition interrupted when we have in-the-head cognitive systems, of which there are, on my own pluralistic view, plenty? Consider one such example: that of *planning for a holiday*, where, let us suppose, all that cognitive activity goes on inside one's head. Could we have a case where that process begins, but it is terminated before completion because the agent, the planner, meets an untimely end? Would this be an example of *planning without a planner*? If there were a way for the activity of planning (at least a little bit of it) to continue beyond the death of the planner—say, through setting planning activity in motion in a self-continuing way—then we would have, I think, a case that ran parallel to the original case of the squashed bug. Suppose that shortly

before meeting an untimely end, the planner were to hand-off sub-tasks to mindless robots, which then continued to generate alternatives, collect information on option details, algorithmically assign overall values to those options, etc. Then it seems to me that we have a case of planning interrupted, where this is properly described as involving ongoing planning activities without there any longer being a planner. Since planning is a cognitive activity, this is an example of cognition without a cognizer, a subject of cognition.

I suspect that one could similarly view a range of other examples—understanding spoken language, smelling a flower, remembering an event—once one elaborates on the functional structures that underwrite those processes. But that will have to remain a suspicion here, for reasons of space.

Finally, the intuition *no cognitive processing without an agent of cognition* could be preserved by someone willing to describe all such cases as involving what one might call *agent shifting*. That is, consider that the process begins with Agent 1, and when Agent 1 goes out of existence, Agent 2 comes into existence as the agent of the continuing process. In the digestion case, Agent 2 would be the part of Agent 1 that continues to exist—the injected prey; in the case of planning (understanding speech, etc.), Agent 2 would be identified as the bots that continue to plan (or process speech) after Agent 1 is terminated.

This view would, in effect, restore the parallel between the biological and cognitive cases by claiming that there really is a biological agent present whenever there is (say) digestion going on, even if it is not the original agent. But ditto for cognition. I leave development of this view for those more adventuresome and promiscuous about agency than am I. My point here is just that such a development is very unlikely to buttress the initial claim that there is an important asymmetry between the cognitive and biological cases that would provide a basis for viewing extended *cognition* as facing some special kind of challenge.

Acknowledgements

This paper is based on discussions that arose during a number of invited talks over the past six years, and I thank participants for the ongoing conversation here: at the International Society for Theory in Psychology Annual Meeting at York University in Toronto in June 2007; at the conference Perception, Action, and Consciousness at the University of Bristol, Bristol in July 2007; at the Cognitive Science Summer Institute in Montreal in July 2008; at The Mind Forum in Helsinki in September 2008; at the University of Colorado, Boulder and at UNAM in Mexico City in March 2009; at the Canadian Atlantic Region Philosophers' Association Meeting in Halifax, October 2010; at the workshop on Embodied, Distributed, and Extended Cognition at the Autonomous University of Barcelona in Barcelona in March 2011; at the Calgary History of Medicine group monthly meeting in Calgary in March 2012; and at the Sante Fe Institute in New Mexico in November 2012. Special thanks to Rob Rupert, Ned Block, Matt Barker, Christian Lee, Carol Cleland, Kathrin Koslicki, and Georg Theiner for more persistent interrogations.

References

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

Burge, T. (1979). Individualism and the mental. In P. French, T. Uehling & H. Wettstein (Eds.), *Midwest studies in philosophy, vol. 4: Metaphysics* (pp. 73–121). Minneapolis, MN: University of Minneapolis Press.

Burge, T. (1986). Individualism and psychology. *Philosophical Review*, *95*, 3–45.

Chomsky, N. (1959). Review of B. F. Skinner's *Verbal behavior*. *Language*, *24*, 163–186.

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

Clark, A. (2008). *Supersizing the mind*. New York: Cambridge University Press.

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

Fodor, J. A. (1987). *Psychosemantics: The problem of meaning in the philosophy of mind*. Cambridge, MA: MIT Press.

Harman, G. (1988). Wide functionalism. In S. Schiffer & D. Steele (Eds.), *Cognition and representation* (pp. 11–18). Boulder, CO: Westview.

Kirsh, D. (2009). Problem solving and situated cognition. In P. Robbins & M. Aydede (Eds.), *The Cambridge handbook of situated cognition* (pp. 264–306). New York: Cambridge University Press.

Levy, N. (2007). *Neuroethics: Philosophical challenges for the twenty-first century*. New York: Cambridge University Press.

Menary, R. (Ed.). (2010). *The extended mind*. Cambridge, MA: MIT Press.

Newell, A., Shaw, J. C., & Simon, H. (1960). Report on a general problem-solving program. In S. de Picciotto (Ed.), *Information processing: Proceedings of the International Conference on Information Processing, UNESCO, Paris 15–20 June 1959* (pp. 256–264). Paris: UNESCO.

Patterson, S. (1991). Individualism and semantic development. *Philosophy of Science*, *58*, 15–35.

Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, *101*, 389–429.

Rupert, R. (2009). *Cognitive systems and the extended mind*. New York: Oxford University Press.

Sawyer, R. K., & Greeno, J. G. (2009). Situativity and learning. In P. Robbins & M. Aydede (Eds.), *The Cambridge handbook of situated cognition* (pp. 347–367). New York: Cambridge University Press.

Segal, G. (1989). Seeing what is not there. *Philosophical Review*, *98*, 189–214.

Shoemaker, S. (1981). Some varieties of functionalism. *Philosophical Topics*, *12*, 93–119.

Sneddon, A. (2011). *Like-Minded: Externalism and moral psychology*. Cambridge, MA: MIT Press.

Theiner, G. (2008). *From extended minds to group minds: Rethinking the boundaries of the mental*. (Doctoral dissertation). Indiana University, Bloomington, Indiana.

Theiner, G., & Wilson, R. A. (2013). Group mind. In B. Kaldis (Ed.), *Encyclopedia of philosophy and the social sciences* (pp. 401–404). Thousand Oaks, CA: Sage.

Turner, S. (2000). *The extended organism: The physiology of animal-built structures*. Cambridge, MA: Harvard University Press.

Wheeler, M. (2010). In defense of extended functionalism. In R. Menary (Ed.), *The extended mind* (pp. 246–270). Cambridge, MA: MIT Press.

Wilson, R. A. (1994a). Causal depth, theoretical appropriateness, and individualism in psychology. *Philosophy of Science*, *61*, 55–75.

Wilson, R. A. (1994b). Wide computationalism. *Mind*, *103*, 351–372.

Wilson, R. A. (1995). *Cartesian psychology and physical minds: Individualism and the sciences of the mind*. New York: Cambridge University Press.

Wilson, R. A. (1999). Philosophy: Introduction. In R. A. Wilson & F. C. Keil (Eds.), *The MIT encyclopedia of the cognitive sciences* (pp. xv–xxxvii). Cambridge, MA: MIT Press.

Wilson, R. A. (2001). Two views of realization. *Philosophical Studies*, *104*, 1–30.

Wilson, R. A. (2004). *Boundaries of the mind: The individual in the fragile sciences: Cognition*. New York: Cambridge University Press.
Wilson, R. A. (2005a). *Genes and the agents of life: The individual in the fragile sciences: Biology*. New York: Cambridge University Press.
Wilson, R. A. (2005b). Collective memory, group minds, & the extended mind thesis. *Cognitive Processing*, *6*, 227–236.
Wilson, R. A. (2007). A puzzle about material constitution and how to solve it: Enriching constitution views in metaphysics. *Philosophers' Imprint*, *7*(5), 1–20.
Wilson, R. A. (2009). The transitivity of material constitution. *Noûs*, *43*, 364–378.
Wilson, R. A. (2010a). Extended vision. In N. Gangopadhyay, M. Madary, & F. Spicer (Eds.), *Perception, action and consciousness* (pp. 277–290). New York: Oxford University Press.
Wilson, R. A. (2010b). Meaning making and the mind of the externalist. In R. Menary (Ed.), *Externalism* (pp. 167–188). Cambridge, MA: MIT Press.
Wilson, R. A. (2010c). Review of Robert Rupert, *Cognitive systems and the extended mind*. *Notre Dame Philosophical Reviews*. Available at http://ndpr.nd.edu/news/24309/?id 19128
Wilson, R. A., & Craver, C. (2007). Realization. In P. Thagard (Ed.), *Handbook of the philosophy of science, vol. 12: Philosophy of psychology and cognitive science* (pp. 81–104). Amsterdam: Elsevier.
Wilson, R. A., & Foglia, L. (2011). Embodied cognition. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2011 ed.). Available at http://www.plato.stanford.edu/entries/embodied-cognition/
Wilson, R. A., & Lenart, B. (forthcoming). Extended mind and identity. In J. Clausen & N. Levy (Eds.), *Handbook of neuroethics*. London: Springer.