

Tilburg University

Embedded Cognition and Mental Causation

Keijzer, F.; Schouten, M.K.D.

Published in:
Synthese

Publication date:
2007

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
Keijzer, F., & Schouten, M. K. D. (2007). Embedded Cognition and Mental Causation: Setting Empirical Bounds on Metaphysics. *Synthese*, 158, 109-125. <http://www.springerlink.com/content/4673957k25p247q1/>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Embedded cognition and mental causation: setting empirical bounds on metaphysics

Fred Keijzer · Maurice Schouten

Received: 6 October 2005 / Accepted: 15 May 2006 / Published online: 26 July 2006
© Springer Science+Business Media B.V. 2006

Abstract In this paper, we challenge Jaegwon Kim's version of neural reductionism according to which the causal powers of mental properties are pre-empted by those of neural properties. Using empirical and theoretical developments from the field of embedded cognition, we articulate and defend a notion of process externalism that extends Clark and Chalmers' notion of an extended mind. We argue that process externalism undermines one of Kim's key premises leading to the alternative conclusion that mental causation cannot be reduced to neural causation. Instead, mental properties have their own new causal powers just like other scientifically established macroproperties.

Keywords Kim · Mental causation · Embedded cognition · Externalism

1 Introduction

Does my desire for milk cause me to walk to the kitchen and to open the fridge, or is my brain the real causal factor? For many years, the status quo within philosophy of mind and cognition consisted of a nonreductionist version of physicalism where either beliefs and desires or at least cognitive processes were held to be causally efficacious factors autonomous from the brain. Nowadays, all this is under attack. Progress in the cognitive neurosciences makes it clear that mind and brain are tied much closer together than we used to think when the multiple realizability argument was fresh (Bickle, 1998, 2003; Churchland, 2002). In its wake, identity theory is on the rise again

F. Keijzer (✉)
Faculty of Philosophy, University of Groningen,
Oude Boteringestraat 52, 9712 GL Groningen,
The Netherlands
e-mail: f.a.keijzer@rug.nl

M. Schouten
Faculty of Philosophy, Tilburg University,
P.O. Box 90153, 5000 LE Tilburg,
The Netherlands
e-mail: m.k.d.schouten@uvt.nl

(Polger, 2004), and within metaphysics, arguments are developed that raise principled problems for the notion of mental causation and make the reduction of mind to brain seemingly inescapable. The work of Kim is particularly influential here (Kim, 1996, 1997, 1998, 2005). He forcefully argues that neural causes pre-empt mental causation: If a neural cause suffices to account for the mental consequence, what role is there to play for an additional mental one? When taken seriously, his metaphysical claims for a new reductionism make the notion of mental causation as something different from neural causation extremely problematical. He sketches a gloomy future for non-reductive notions of mind, for instance those based on supervenience. As he notes at one place: “all roads branching out of physicalism may in the end seem to converge at the same point, the irreality of the mental” (1998, p. 119).

Kim’s argument has huge implications for philosophy of mind and cognition, as well as psychology and cognitive science, as it puts doubt on the future possibility of any psychology or cognitive science that is not ultimately a form of neuroscience. In the light of these serious philosophical and empirical implications, it is not surprising that these claims have not gone uncontested (Block, 2003; De Muijnck, 2003; Gozzano, 2003; Loewer, 2001; Noordhof, 1999). Ross and Spurrett (2004) have now even written a “defense manual for cognitive and behavioral scientists” against skeptical metaphysicians. Most of these responses to Kim’s analysis focus on internal difficulties with Kim’s argumentation and challenge the validity of his conclusions in this way. However, in the present paper, we propose a different approach that does not contest the backbone of Kim’s metaphysics. Instead, we will focus on a key assumption that plays a crucial role in Kim’s metaphysical argument against the reality of mental properties: Mental and neural properties are instantiated by the same physical entity. We will argue that this particular assumption is questionable as it is considered false, on both philosophical and empirical grounds, in an increasingly wide circle of cognitive scientists. Ironically, when this particular assumption of Kim is modified on the basis of these new ideas, Kim’s argument—left otherwise intact—enables the conclusion that mental causation is real and separate from neural causation.

The philosophical and empirical work on cognition that has these interesting implications for Kim’s analysis comes from *embedded cognition*. This is a general label that refers to work carried out in fields like developmental psychology, robotics, neuroethology, movement science, philosophy, and consciousness studies.¹ Embedded cognition ties mental phenomena closely to bodily and environmental factors. For our present purpose, the key implication of embedded cognition is that it leads to a form of externalism concerning psychological processes: Mental processes are not only constituted by processes within our nervous system, but *also* incorporate bodily and environmental processes. We will call this view *process externalism*. Process externalism must be contrasted with Putnam’s and Burge’s standard meaning externalism. It is much more akin with Clark and Chalmers’ active externalism (Clark, forthcoming; Clark & Chalmers, 1998) and even more so with Hurley’s vehicle externalism (Hurley, 1998a, 1998b, 2001). Process externalism goes directly against Kim’s claim that psychology is ultimately a matter of neuroscience alone. Note that the idea of process externalism and its kin has raised a lot of criticism (Adams & Aizawa, forthcoming; Grush, 2003; Rupert, 2004; Wilson, 2002). However, we want to stress from the outset

¹ Some of the most important references are for developmental psychology Thelen and Smith, 1995; for robotics Brooks, 1999; Pfeifer and Scheier, 2001; for neuroethology Beer, 1995; for philosophy Clark, 1997; Hurley, 1998a, 2001; Van Gelder, 1995; and for consciousness studies O’Regan and Noë, 2001.

that process externalism actually constitutes a relatively mild claim, which states only that not all mental processes can be exhaustively described by neural processes, or processes occurring within the skin.

By bringing Kim's argument against mental causation and embedded cognition together, we aim to do two things. First, we hope to counter Kim's metaphysical claims concerning the irreality (a) of non-neural mental causation and (b) of any future psychology which is not merely a branch of neuroscience. Second, we will try to show more generally that his metaphysical conclusions are actually held hostage to the particular outcomes of empirical studies of the nature of psychological processes.

The paper has the following structure. First, we briefly review the basic structure of the causal exclusion problem for mental causation and Kim's position on this issue. Second, we will articulate the notion of process externalism as it arises within the context of embedded cognition and go over some empirical evidence which provides, in our eyes, a convincing case for the limited claims of process externalism. Third, we will sketch how process externalism can give rise to a widely constructed form of mental causation. Finally, we will conclude that Kim's causal exclusion argument, when combined with process externalism, provides a solid case for real—that is more than strictly neural—mental causation.

2 The exclusion of mental causes

Physicalism is a fundamental constraint that philosophy of mind has to take into account. How does this constraint bear on the ideas that our thoughts are the causes of our actions, like when I open the refrigerator door because I have a *desire* for milk and I *believe* that I can find some in the fridge? For a long time, a nonreductive consensus held that mental causation could be reconciled with the strictures of physicalism.

On Kim's analysis, attempts at nonreductive solutions like supervenience cannot save mental causation because of the exclusion principle: If supervenience is accepted, then there will inevitably be a problem with the causal efficacy of supervenient properties. Any causal story involving beliefs and desires will be pre-empted by a more basic story in terms of the neurophysiological processes. My desire for milk *M* caused my physical property *P**, but this action was also causally necessitated by physical property *P* (e.g., the firing of hypothalamic neurons). Given supervenience, *M* is causally active because *P* is, so the question is: How can *M* really be a cause of *P**? Kim concludes that it cannot and that nonreductive versions of physicalism now face hard times. The constraints set by physicalism cannot be reconciled with mind as an autonomous causal factor in addition to the brain. Kim sees reductive physicalism as the best we can get from physicalism, while mental causation can be nothing more than neural causation.

We are not convinced that Kim's analysis of nonreductive physicalism and mental causality is correct. However, our point will be that the causal efficacy of the mental can be maintained *within* the confines of Kim's own "robust physicalism" (1998, p. 120). The groundwork for this point comes from his treatment of the generalization argument, which is often brought up against the exclusion principle (e.g., Baker, 1993; Burge, 1993; Fodor, 1991). The generalization argument states that the exclusion principle commits Kim to the claim that no macrocausation is possible at all. The exclusion principle would apply to all forms of higher-level causation. Chemical and biological

forms of causation would also be pre-empted by more basic lower-level ones until one reaches a bottom physical level and only there true causal powers would reside. As this conclusion is thought to be absurd, the implication is that the exclusion principle cannot spell real problems for mental causation.

Kim's opinion goes in the other direction: Exclusion stands, and there is a problem with explaining why all macrolevel causation does not fall victim. In reply to the latter problem, he states that physicalism does not imply microphysicalism. New causal powers can be had for micro-based macroproperties, which "*can, and in general do, have their own causal powers, powers that go beyond the causal powers of their micro-constituents*" (Kim, 1998, p. 85). Because, for instance, biological properties often apply to larger, macroscopic, entities compared to chemical properties, the former can have new causal powers relative to the latter. However, according to Kim, this way of having macrocausation is not open for mental causation. There is a fundamental asymmetry between psychology's relation to neuroscience and, for instance, biology's relation to chemistry. In contrast to the latter, mental and neural properties are properties of the same object and reside at the same level: "they are both had by human beings and other sentient creatures" (Kim, 1998, p. 117). Hence, the argument just given does not apply here. Mental properties are second-order properties defined over their first-order physical (neural) realizers in much the same way as 'primary color' is a second-order property relative to the domain of colours. As the causal powers of second-order properties are had in virtue of those of the first-order base properties, the latter pre-empt the causal powers of the former. As Kim puts it: "the causal role of a mental property had by me is threatened with pre-emption by another property, a neural property, also had by me" (Kim, 1998, p. 117). Both properties have identical causal powers. Hence, the generalization argument will not save mental causation, says Kim.²

In the following, we will work with Kim's analysis without questioning its general structure. What we will do is focus on one important assumption in his argument. According to Kim, mental and neural properties apply to the same entity, as both are had by human beings. We will question this assumption and argue that both sets of properties apply to different entities at different levels of organization. The issue is highly relevant. If mental properties can be ascribed to a mental entity constituted by psychological processes that incorporate brain processes as a mere part, then Kim's way of saving macrolevel causation—or similar ways of saving macrolevel causation—would also apply to mental causation. The conclusion that can then be drawn from Kim's analysis is that mental causation is safe and sound, with its own, new causal powers. In this respect, our approach comes very close to Pereboom's (2002)

² We based this analysis primarily on Kim's 1998 book. In his recent book *Physicalism, or something near enough*, Kim does not return to the level-order distinction. He also seems, implicitly, now to be more equivocal concerning 'new' causal powers for macroproperties (2005, pp. 57–69). However, his main assumptions on which we will build our argument remain intact. Kim still claims that "the supervenience argument does not eliminate all macrocausation" (2005, p. 56). His example here is that a ball breaks a window, not the ball's microconstituents. Kim also becomes highly explicit concerning the close, reductive link between the mental and the neural, claiming that "we already know the neurophysiology involved well enough" to conclude that mental causation is just neural causation (2005, p. 153). This last issue is crucial for our argument below as we go against such an identity claim on empirical grounds. If Kim now has second thoughts about macrocausation in general, our point that mental causation is no worse off than any other form of macrocausation still stands. In addition, mental causation can ride with any rescue operation that aims to save the general notion of macrocausation from any of the problems that Kim discusses.

‘robust nonreductive materialism’ and can be seen as a way to make this nonreductive position more concrete.³

Before proceeding to embedded cognition, it is important to have a look at Kim’s assertion that mental and neural properties are had by the same entity. In general, Kim seems to hold it as self-evident that the brain is the exclusive organ of the mind, so much so that he only mentions it in passing (e.g. 1998, p.117).⁴ However, this is not self-evident at all. For example, it is a controversial issue in the philosophy of neuroscience whether psychological and neural explanations are at the same level or not (McCauley, 2001; Churchland & Churchland, 2001). In addition, if Kim is to be right he must mean a very particular kind of neurobiological property—and not, for example, a property like being serotonergic. While we could say that we ‘feel serotonergic’, meaning that we are happy or so, we would not be actually serotonergic in the sense as it applies to certain neurons: secreting and being sensitive to serotonin. Some of my neurons are sensitive to serotonin, but I am not. The general point has been argued at length by Bennett and Hacker (2003). Finally, there are many different kinds of neurobiological properties that can apply to various levels of brain organization. Being serotonergic and being frontal apply to different levels of neural organization and not to the same entity. Given these considerations, Kim ought to—but does not—provide an account that specifies which neurobiological properties he has in mind such that they are had by me in the same way as my mental properties are had by me. The example that Kim offers to illustrate his point is not helpful here. He states that: “Socrates’ goodness supervenes on his honesty, generosity courage, and wisdom, and it is the same person, Socrates, who instantiates both these subvenient virtues and supervenient goodness” (1998, p. 86). However, in this example there are no neurobiological properties involved, only psychological ones, which apply to Socrates, not to his brain. Kim’s example does not strengthen or clarify the idea that psychological and neurobiological properties apply to the same entity. To conclude, Kim’s claim that mental and neurobiological properties apply to the same entity is a questionable assumption, which he does not adequately defend. In the following, we will argue that this claim comes easily apart when current developments in the field of embedded cognition are brought to bear on this issue.

3 Embedded cognition and process externalism

Are mental properties and neural properties instantiated by the *same* entity? Disregarding the problems involved in making this general idea sufficiently precise, a strong reason for thinking that this question must be answered in the affirmative comes from a widely held view in both philosophy of mind and cognitive science: mental processes are physically constituted by brain processes. If you have the brain, you will also have the mind. This general idea will here be referred to as process *internalism*.

In contrast, the field of embedded cognition involves empirical and theoretical work which stresses the intricate and essential connections between brain processes, bodily processes and environmental processes.⁵ The conceptual claim here, inspired

³ We thank Carl Gillett for drawing our attention to the link with Pereboom’s work.

⁴ As said in footnote 2, he has become more explicit about this point in his 2005 book.

⁵ See for instance Beer (1995), Brooks (1999), Clark (1997), Hurley (1998a, 2001), Keijzer (2001), Noë (2004), O’Regan and Noë (2004), Pfeifer and Scheier (2001), Thelen and Smith (1995), and Van Gelder (1995).

and backed by empirical evidence, is that these connections are so extensive and important that one must conclude that at least some psychological processes cannot be localized in the brain alone, but must be interpreted as spread out across the body and the environment. Embedded cognition states that it is misleading to suggest that body and environment are simply add-ons for autonomously operating mental or brain processes. In contrast, embedded cognition leads to what we will call process *externalism*: some mental processes are not fully constituted by internal brain, or even bodily processes, but also involve environmental processes. A subset of mental processes, in this view, are literally co-constituted by environmental processes.

At this point, our reading of the mainstream position as process internalism might be questioned: Nobody did ever believe that a brain on its own would could realize mental processes. It is a general and obvious background assumption that a brain must be connected to a body and an environment to operate normally. The claim that brain processes physically realize mental processes does not imply that body and environment are not also important. In response, we must stress that embedded cognition involves a stronger claim than just 'being important.' The standard view, process internalism, does acknowledge that bodily and environmental factors are essential as boundary conditions for brain processes. However, process externalism emphasizes that certain aspects of these so-called boundary conditions are literally a part of (at least some) mental processes. This is a radical claim that differs significantly from previous views in cognitive science and philosophy of mind.⁶

At this point, we want to stress that the introduction of process externalism provides the key issue in our case against Kim's conclusions concerning mental causation. Process externalism goes directly against the idea that mental causation can be no more than neural causation. We also want to stress that process externalism may be controversial—however see our relatively modest elaboration below—but that a single counterexample of an extended mental process suffices to break the identity claim on which Kim build his case that mental causation is pre-empted by neural causation. Kim's position is thus very feeble to start with.

Probably the best known, and most widely discussed, attempt to articulate what we call process externalism comes from Clark and Chalmers (1998, see also Clark, forthcoming). They introduced the phrase *active externalism* which they distinguished from Putnam and Burge's meaning externalism. Meaning externalism is compatible with the idea that ongoing thought processes can be exhaustively described in terms of processes occurring within the brain (Hurley, 1998a). Active externalism claims that thought processes are not strictly tied to the brain. Clark and Chalmers' version of process externalism focuses on the use of artifacts as mind extensions. For example, my mobile phone can, under certain conditions, be interpreted as part of my memory, even though it is external to my body. We are sympathetic to Clark and Chalmers' ideas. However, we think that process externalism must,⁷ and can, be cast in a stronger way, by focusing on more basic agent-environment interactions such as those occurring in perception and action more generally. For example, Hurley stresses the importance of continuous dynamic feedback loops through the environment as

⁶ The exact implications and relations of process externalism with standard meaning externalism still need to be worked out. See Hurley (forthcoming) for a discussion of these relations.

⁷ In particular their choice of memory seems vulnerable to a different interpretation (e.g., Rupert, 2004). One can easily maintain that the differences between externally mediated and strictly brain-based forms of memory are sufficient to reserve the term 'belief' for the latter.

a reason for doubting any strict internalist interpretation of the vehicles that ‘carry’ meaning at a subpersonal level (Hurley, 1998a). The phrase *vehicle externalism* is generally applied here (Hurley, 1998a, 1998b, 2001; Metzinger, 2003; Rowlands, 2003). Vehicle externalism comes very close to our own process externalism. However, as we want to avoid any commitment concerning the personal-subpersonal distinction here, we prefer the phrase process externalism, which is neutral in this respect.

As said, process externalism amounts to the claim that the linkages between an agent and her environment are so important to her ongoing functioning as an agent that it is necessary to take relevant environmental features as belonging to the set of processes that together constitute mentality. Versions of process externalism are currently hotly debated.⁸ As the focus of the present paper is on Kim’s analysis of mental causation, we will not provide an elaborate discussion and defense of process externalism here but simply refer to this ongoing work.⁹ Nevertheless, to give a general idea of the issues involved, we will say something about what Adams and Aizawa (forthcoming) call the causal-constitution error, and Hurley (forthcoming) the ‘causal-constitution error’ error.

Adams and Aizawa argue that what we call process externalism involves making a causal-constitution error. Process externalism mixes up constitution and causal relations. In their view, an agent is physically constituted by the matter that makes up her body and this is different from the causal relations between this body and the external environment. These causal relations are granted to be essential for the upkeep of the body and brain—just think of the role of the environment as a supplier of oxygen—but this does not make it in any way a part of the bodily system. Hurley (forthcoming), on the other hand, claims that the idea that there is a causal-constitution error is itself an error, namely the error of assuming *without argument* that external portions of psychological explanations are causal rather than constitutive. We agree with Hurley that it cannot simply be assumed that process externalism makes a causal-constitution error. It will have to be decided by empirical evidence and theoretical developments where we are going to draw the boundaries around the mental.

Concerning theoretical developments, we want to say the following. Embedded cognition involves an important reconstruction of the mind-concept, founding it in perception-action capabilities and treating inner thought as an evolutionary and developmentally late extension of those capabilities. Fish mentality arose before human minds and there is no adult human mind which hasn’t been a baby mind first. Embedded cognition involves a naturalistic conceptualization of mental processes which starts out from its evolutionary and behavioral context, rather than naturalizing the existing notion of a rational, reflective mind, which is ultimately based on our own human, adult experience. The key implication here is that *the fundamental criteria for mentality do not consist only of intentionality and consciousness but also of perception-action relations*. Externalist claims concerning the constitution of psychological process are themselves a reflection of a conceptual extension concerning the core

⁸ For defenses of the claim that mental processing extends into the environment see, for example Clark and Chalmers (1998), Clark (forthcoming), Haugeland (1998), Hurley (1998a, 1998b, 2001), and Rowlands (2003). For an additional defence and overview of different externalist claims see Hurley (forthcoming). For criticisms of externalist thinking see, for example Adams and Aizawa (forthcoming), Grush (2004), Rupert (2004), and Wilson (2002).

⁹ We will give our own elaboration and defense of process externalism with respect to the causal-constitution problem, the brain in a vat, the unplugging problem and other issues elsewhere (Keijzer & Schouten, in preparation).

issues of mental processes. Given the importance of ongoing dynamical couplings with the environment for intelligence, these couplings can become themselves an important indicator of mentality (Van Gelder, 1995), and precisely these dynamical couplings make it theoretically plausible to start thinking in terms of mental processes as dynamically extended.

In the light of this theoretical change in ideas concerning what makes up mentality, we think process externalism becomes relatively easy to establish. In the following we will discuss some of the findings regarding visual perception that have most generally led to conclusions similar to our process externalism. Most notable among these is the phenomenon of change blindness (Rensink, 2002; Simons & Levin, 1997), which has also been used for similar purposes by O'Regan and Noë (2001; see also Noë, 2004). Experimental findings show that large changes in a visual scene, such as objects changing colors or appearing and disappearing objects, often go unnoticed as long as the actual change is masked and the person is not paying attention to the changed feature (Simons, 2000). The effect can be highly remarkable. In one experiment, subjects did not even notice that the person they were talking to face to face was replaced by someone else (Simons & Levin, 1998). Change blindness provides a refutation of the subjectively plausible idea that seeing is like making a photograph. A camera opens its shutter for a very short period, then closes again, leaving a complete trace of the light pattern on its inner photographic film. Vision seems similar in that we (and our brains) seem to take in a whole visual scene with a single momentary glance, after which we have acquired a subjectively accessible internal representation of the external scene. Change blindness shows that it is highly questionable that we ever acquire such rich internal representations of the visual field. If we had, we would notice the changes. In contrast, visual perception has to be interpreted as an ongoing process where the visual scene is scanned according to need, and at no particular moment is there a full internal representation of the external scene present inside the brain (Churchland, Ramachandran & Sejnowski, 1994). Our subjective experience of having access to all the details within a wide visual field derives from our capability of scanning this scene at will. In other words, visual perception does not take place inside the brain, but is realized by the ongoing interactions between brain, body and the external visual scene.

We think the phenomenon of change blindness is very convincing evidence for our case. Still the interpretation of these experiments is not yet unanimous (Simons, 2000). However, with hindsight, the same conclusion can also be drawn from the basic textbook case of visual saccades.¹⁰ Our eyes differ from a camera as only a very small spot in the center of our retina is very sensitive to visual details and to color, while the rest of the retina generates a much coarser image. To compensate, normal vision involves four to five small eye movements, saccades, a second that scan the visual field continuously. Only by and during such an active and ongoing process do we acquire the general detailed views of normal vision. When I see the park outside, my eyes saccade across this scene in order for me to see a definite image consisting of trees, grass, a bridge, and houses all around. There is never at any given time a strictly internal state that corresponds with my consciously seeing this scene. To see the park is an ongoing active process, involving saccades and the external scene as necessary ingredients. As Ballard, Hayhoe, Pook and Rao (1997) note, eye movements are the

¹⁰ This example has also been used by O'Regan and Noë (2001) for their much stronger claim concerning visual consciousness.

best example of a fast sensory-motor primitive. Thus, while seeing seems intuitively a process occurring within the brain as the internal effect of an external input, it is physically constituted by the close coupling between organism and environment.

Before rounding off we will introduce an analogy that might help to make our point more intuitively plausible. How do airplanes fly? The rough story is that the forward speed of the airplane creates a pattern of air movement around the wings, which generates a higher air pressure on the underside of the wings, and a lower pressure on top of the wings. Together this creates an upward force which lifts the plane off the ground into the air. Our intuitions here are that flying consists of a process which involves both aspects of the plane itself and the dynamics of the airflow around it. Both are necessary ingredients of this form of flying, in contrast to say ballooning or rocketing. Airplane flight has to be interpreted in an externalist way, akin to what we call process externalism for psychological processes. Obviously, there is no claim that the airplane itself somehow extends into the environment. The claim is only that the processes which keep the airplane above the ground include ingredients external to the plane. The same goes for vision, we hold. Vision is a continuous and ongoing process, involving both bodily activity and perceptual elements and the environmental link between the two. It remains very controversial whether this implies that our *experience* has to be interpreted in an externalist way (e.g. Hurley, forthcoming; O'Regan & Noë, 2001; Noë, 2004). However, we think that vision more widely construed as the ongoing coordination between an organism and its environment is both an example of a clear and uncontroversial psychological process, and a case where otherwise physically separate ingredients—internal and external to the body—become bound together into a single psychological process of seeing.

In sum, process externalism is a claim that can be derived from embedded cognition: at least some psychological processes are co-constituted by neural, bodily and environmental processes. As a general claim concerning mental processes, process externalism remains disputed. However, we find the case of visual perception, as sketched above, very convincing evidence for process externalism with respect to this particular psychological process. How far this conclusion will be extended toward other psychological processes remains to be seen. Embedded cognition is now making an important impact on the cognitive sciences as strictly internalist views are giving way in a certain degree to more externalist interpretations (Pecher & Zwaan, 2005). However, at present it is much too early to tell how this debate will develop. Nevertheless, this general inconclusiveness concerning process externalism in other psychological areas does not affect our present case against Kim's analysis. For his claim that mental causation is in deep trouble, he requires the very strong claim that mind and brain (or another non-mental physical entity) are identical. Mental and physical properties must apply to the same entity. In comparison, our claim concerning process externalism is very modest. To defeat Kim's *identity* claim, a single counterexample of an extended psychological process will suffice. Visual perception provides such a counterexample and all on its own is sufficient to undercut Kim's analysis concerning mental causation.

4 Process externalism and mental causation

What we hope to have shown so far is that some mental processes are co-constituted by neural, bodily and environmental processes. We have assumed further that mental

processes together constitute the mind. This implies that mental properties and neural properties do not apply to the same entity and that therefore mental properties cannot generally be cast as second-order properties over first-order neural realizers. Working with Kim's (1998) analysis of macroproperties, process externalism can now be cast as a way of gaining new, and particularly mental, causal powers for mental properties. Kim's own analysis and defense of macrocausation for chemical and biological domains will also apply to the psychological domain. Like chemical and biological properties, mental properties occupy a distinct level in a micro-macro hierarchy and in virtue of this fact it can be claimed, following Kim, that mental properties have their own causal properties over and above non-psychological first-order properties.

Strictly speaking, we should be finished here, our aim to overcome Kim's denial of mental causation achieved by our defense of process externalism. Nevertheless, it will be beneficial to strengthen our case further in a number of ways. First, we have been equivocal in our use of the words mental and psychological. We should give some clarification, in particular when it comes to a defense of mental causation. Second, we should say something about the problematical relation between externalism and mental causation, which makes it strange to use it in a defense of mental causation. Third, we want to give an example of this extended mental causation, and fourth, give an additional argument, derived from the *screening off* literature, to defend the causal relevance of those higher-level psychological entities with respect to their lower-level constituents.

So far, we have been vague with respect to what we take the mental to be, sometimes referring to the beliefs and desires of folk psychology and sometimes to cognitive science or scientific psychology. Mental causation is usually placed in a folk psychological context where it is phrased as the question how I, as a mental entity, can influence the physical world. In this view, mental causation easily becomes a mysterious influence on neural and bodily processes. We want to stress that embedded cognition is not intrinsically tied to folk psychology but to scientific psychology and cognitive science. It is in this context that we want to position and develop a notion of mental causation, which is not necessarily tied to the folk psychological view that my desires must be cast as the cause of my actions. In this scientific context, the question concerning mental causation relates to the issues whether or not the psychological system or entity as a whole has its own, new, causal powers that are not present in the separate components of the entity. From such a perspective, there is nothing intrinsically mysterious to mental causation. In addition, and as said before, given embedded cognition, the criteria for mentality are not restricted to intentionality and consciousness but also involve perception-action relations. It is the latter aspect of embedded cognition, which makes it fairly easy to cast psychological entities as environmentally extended, we hold.

A second issue concerns the question how a psychological entity that extends into the environment can have causal powers. Traditionally, meaning externalism has been causing problems for the notion of mental causation. However, process externalism involves a very different claim, which does not have these problems (Hurley, forthcoming). Now, although we don't have the same problems as a meaning externalist, we may be said to have others. Someone might deny that causality must, or even can be ascribed to any extended mental entity. The problem consists of an incongruity between the psychological and, say, the biological case that will undermine the causal efficacy of the mental. One can argue that the mental, but not the biological, extends beyond a thing's intrinsic microstructure and therefore the mental, but not the biological, violates what many take to be a fundamental constraint on scientific theorizing:

individualism. Individualism is the thesis that causal powers must supervene on local, intrinsic microstructure and this, Fodor adds, “is a constitutive principle of science” (1987, p. 45). Extended mental causation involves spooky action at a distance and must be prohibited from entering scientific theories, or so the objection could go.

In our view, this objection misconstrues process externalism. The incongruity between the mental and, say the biological, is not there. According to process externalism, psychological processes simply have an extended microstructure, bound together by ongoing perception-action couplings. In this, psychology is not intrinsically different from biology. Biological processes are constituted by a host of chemical and physical processes that tie themselves together to form biological entities with characteristics like metabolism and morphogenesis, characteristics not had by these constituents. Chemical interactions take place on molecular scales and from this small-scaled perspective it might seem that a cell or organism is an extended entity that encompasses many different processes and structures often as far apart as millimeters and even meters. What from a chemical perspective would be a hazardous collection of different and far-flung constituents is brought together into a single living entity, which can act as a local, though comparatively large, intrinsic microstructure.

The biological example makes clear that ‘local’ has to be interpreted in the context of the entities that one discusses. For psychology this would imply that one can have extended psychological processes which nevertheless do not contradict the locality constraint. As long as ongoing perception-action couplings tie the different components into a unit, they can be part of a local, even though extended, intrinsic microstructure.

Of course, in comparison to living systems these extended psychological entities do not have a clear and bounded material constitution. However, we can refer back to dynamical theories within embedded cognition which precisely stress the point that the material boundary provided by the skin does not necessarily provide a suitable boundary for the dynamical patterns that arise in perception-action coupling (Beer, 1995, 2000; Van Gelder, 1995). From a dynamical perspective, one should talk about trajectories that are shaped both by forces inside and outside of the physical body. In more mundane terms, and going back to our earlier example of flying, the intricate interactions between air and wing together constitute a system of air vortices supporting the plane, and these, together with the plane itself constitute a flying plane. The air vortices may be fleeting but they must be there all the time to make the airplane move through the air. If you want to explain the movement of the plane, the dynamics of both the plane itself and the air moving around it must be taken into account. Flying, but not the airplane itself, involves both internal and external ingredients. The external and plane components become tied together in this singular dynamical phenomenon. A similar story, applies in the psychological case where ongoing, dynamical interactions between agent and environment constitute an agent’s visual perception, which allows him to orient himself with respect to his environment.

How does all of this translate into an example of mental causation? The first thing to acknowledge is that standard examples like how my desire for milk leads me to walk to the fridge are probably not the most suitable to be cast in a process externalist form. That is, as long as one focuses on the initiating desire, which should set my body into motion. However, as soon as the action is initiated and I am on my way to the fridge for that glass of milk, my action is an ongoing affair that involves the ongoing scanning of my visual environment, using the results to adjust my movements and so on. Environmental features are continuously picked up by my eyes, which enable

me to direct and redirect my movements in a continuous way, and which eventually lead to me grasping a bottle of milk and pouring it into a glass. For my action to unfold in the way that it does, continuous dynamical feedback is a necessary part of the process that ends with a glass filled with milk. This feedback does not have to be visual, as I may search for the milk in the dark, but in that case tactile stimuli will act as a replacement. Mental causation in this analysis is not something which imposes itself from the brain outward, but is both present in, and resulting from, ongoing actions which include neural, bodily and environmental factors. The effects of such mental causes can themselves again become part of a more extended mental process (Dretske, 1988). For example, getting a glass of milk will involve first moving over to the fridge and then pouring a glass of milk. It is crucial to realize that mental causation in this reading does not aim to replace straightforward physical causes. The physical displacement of the milk from the bottle into the glass is a definite physical process, which can be described as such. The mental cause applies to the larger-scaled perception-action event that leads to a situation in which the milk was poured from the bottle into the glass rather than not. Without any ongoing perception-action process, the milk would have stayed put.

To make a more general case for our way of casting extended psychological processes as a solid causal factor we will round off with a short discussion of the so-called *screening off* criterion as a way to allocate causal responsibility across micro-macro levels (see also Schouten & Looren de Jong, 2004).¹¹ A property M screens off property P if adding P does not improve the prediction or explanation of outcome R, whereas M does improve the prediction or explanation of R. That is, M screens off P from R iff the assumption of M makes P statistically and/or causally irrelevant with respect to R, but not vice versa (McClamrock, 1995, p. 45). To put the idea slightly more precise we can put this in terms of conditional probabilities:

$$\begin{aligned} & \text{M screens off P from R iff} \\ & \Pr(R|P \ \& \ M) = \Pr(R|M) \neq \Pr(R|P) \end{aligned}$$

Using the Screening-Off Rule (henceforth, SOR), we can, for example, say that barometer readings are screened off from thunderstorms by atmospheric pressures. As soon as one fixes atmospheric pressure, variations in the height of the column of mercury in a barometer will not further correlate with variations in the probability of a storm. Hence, it is the low air pressure, not the imminent storm, which is causally responsible for the dip in the barometer's mercury column. Thus, screening off allows one to say that more proximal causes screen off more distal ones from their common effects. Now, as this use of screening off applies within a particular level of organization, the question is can we use screening off also to defend macrocausation?

Fortunately, screening off can also be plausibly used in the inter-level case. One may counterfactually fix the higher-level, mental properties and in whatever way the underlying microphysical properties will vary (given implementational plasticity, see Wilson, 1995, p. 194), the behavioral effect will be fixed as well. Hence, behavior can be explained in terms of mental processing, screening off the underlying implementational, microphysical processes on which it is dependent. Although cognition is determined by microphysical properties, applying SOR to mental properties, will make them come out as the principal difference-makers with respect to certain types of behavior. More importantly still, for our present case, is McClamrock's claim

¹¹ We thank an anonymous reviewer for raising this issue.

that SOR “can be used even when the property at the ‘preferred’ level of organization is *more distal* than those at other levels of organization” (p. 52). This may be surprising given that SOR typically indicates, as we noted above, causal nearness: distal causes are rendered invisible by proximal causes. A nice example offered by McClamrock to support his claim that distal causes may enter the causal picture is animal camouflage (see also Wilson, 1995, p. 193). Camouflage works, i.e. enhances the organism’s chances of survival and reproduction, as long as the organism remains indiscernible from its surrounding environment. As McClamrock makes clear, there are multiple causal, lower-level causal pathways to the same reproductive success, but the details of the specific lower-level causal route don’t add to the explanation of why the animal manages to escape its predators. When the match/mismatch between organism and surroundings is fixed, its chances of survival will be fixed, in whatever way the lower-level, proximal causal pathways may change (here one may, for instance, think of detailed changes in an animal’s patterns of striping, coloring, etc., or changes in the animal’s underlying physiology or specific changes in the local environment). This provides a telling case in which proximal causes are screened off by distal causes, a case that can also be applied to extended mental causation.

Our example of ongoing perception-action coupling in vision is analogous to McClamrock’s camouflage case: fix the extended processes and in whatever way the microphysical pathways underlying these perception-action couplings may vary, the behavioral effect too will be fixed. Details about the local, microstructural processes in which these perception-action processes are or can be instantiated do not improve our causal explanations of what we are interested in here: how it comes about that I end up quenching my thirst by drinking milk. The phenomenon to be explained correlates with the distal causes regardless of the specific microphysical pathway in virtue of which the distal causes are linked to the behavioral effect (McClamrock, 1995, p. 159; see also Looren de Jong, 2003). In cases like these, SOR (in conjunction with implementational plasticity) makes it clear that causal primacy must be given, not to the underlying neural or microphysical events, but to the extended psychological (perception-action) processes taking place: the latter not the former are causally responsible for the event of me pouring myself a drink.¹²

We conclude that—partially extended—psychological processes are the relevant causal factors for the things we do. At least for certain explanatory purposes, the extended, psychological story is the best causal story to be told. Descending to the microphysical level would mean drowning in a cornucopia of grubby details. Hence, we would like to cast the position we develop here as what Pereboom calls a robust *nonreductive* materialism (see also Schouten & Looren de Jong, 2004). Pereboom states that the nonreductivist “might endorse intrinsic mental properties that instantiate specifically mental causal powers” (2002, p. 517) and argues for the presence of “structure types which cannot be classified as specifically neural, but which must rather be categorized as mental, and which would be intrinsic properties of mental

¹² Another nice example was offered recently in this journal by Chemero and Heyser (2005). They show that in object exploration studies with rodents, recognition memory is causally explained not by reference to neural or molecular processes but by the affordances of the objects explored by the rodents. Here too neural factors are screened off and non-neural factors are causally primary. In fact, they argue for an ecological approach to object exploration, “an approach to cognition in which the object of study is spatially larger than the brain” (p. 419). Their conclusion is entirely in line with our plea for embedded cognition and extended mental causation. We thank an anonymous reviewer for bringing this research to our attention.

states” (p. 518). We see embedded cognition and in particular the example of vision as a way of cashing in such intrinsic mental properties. However, we must stress that even under a reductive interpretation, mental causation would remain specifically mental, analogous to the causal properties of other macroscopic entities.

This embedded view of mental causation diverges from the more standard view in which mental causation refers at heart to causes that derive from an inner mind, whether cast in terms of conscious will, a cognitive decision process or even a brain process. However, the divergence is not a problem for our argument. We have been developing the notion of process externalism in the context of recent empirical developments that go under the label of embedded cognition. Within this context, it makes sense to talk about mental causation as just described. For an empirical case against Kim’s exclusion argument for mental causation, we do not need to abide by mental causation as it is traditionally conceived. Kim claims to make a metaphysical point that is perfectly general and applicable to mental causation without further amendments. When embedded cognition changes our appreciation of the constitution of mental processes, leading to the conclusion that some forms of mental causation derive from an environmentally extended process, then the generality of Kim’s causal exclusion problem for mental causation is directly undercut.

Ironically, in our reading, mental causation might become rather widespread, involving many animals and maybe some machines. How widespread will have to be decided by further empirical and theoretical analysis of psychological processes and the development of suitable criteria that might distinguish visual perception from more basic feedback systems, like those in cruise missiles. Given the situation, it probably is appropriate to talk about psychological causation in this context instead of mental causation. Still, the message remains the same. Process externalism leads to specifically *psychological* causal powers within physicalism.

5 Concluding remarks

We started out with the observation that neural reductionism and new versions of identity theory are on the rise within philosophy of mind. To make this more specific, we focused on an influential argument of Kim, which led him to the conclusion that reductionism concerning the mental is unavoidable and mental causation can be nothing more than neural causation. Our response was to leave the general structure of the argument intact and focus on an essential empirical assumption on which Kim relies: Mental properties and neural properties apply to the same physical entity. By bringing in current developments on embedded cognition, and in particular building up a case for process externalism, we have challenged this assumption. Instead, we claimed that at least some mental processes are constituted by a combination of neural, bodily and environmental component processes, with vision being our primary example. Assuming further that mental processes together constitute the mind, then mental properties and neural properties do not apply to the same entity, and thus mental properties cannot generally be cast as second-order properties over first-order neural realizers.

As a consequence, we claim that Kim’s own analysis and defense of macrocausation for chemical and biological domains will also apply to the psychological domain: mental properties have their own causal properties and are not preempted by other non-psychological first-order properties. They are an important aspect of the natural

world, just like chemical and biological phenomena. Thus, by bringing process externalism to bear on Kim's analysis, the latter does no longer support the conclusion concerning "the irreality of the mental" (1998, p. 119) but can be turned into a strong physicalist defense of mental causation.

Having made our general point, we will round off by discussing two possible ways for Kim to avoid the conclusion we draw here. We will first discuss the relatively preliminary status of process externalism as defended here and the implications for the mental causation issue. Secondly, we discuss the counterargument that micro-based properties are themselves physical and not mental.

Is it possible for Kim to avoid the above conclusions by denying the perspective provided by embedded cognition? After all, the linchpin of our argument consists of the acceptance of a wide constitution for mentality, process externalism, which is not generally accepted so far. Interestingly, even when process externalism is considered controversial, this does not detract from our case against Kim. The very fact suffices that the process externalism issue is now being raised and debated (see footnote 5). Kim claims that the problems concerning mental causation are deep and fundamental. But, his argument is held victim to an ongoing debate on empirical and theoretical issues concerning the constitutive basis of mental processes. Being dependent on such a disputed foundation is in itself damning enough for any metaphysical claim. This point alone suffices as a general lesson that metaphysicists ought not to neglect the empirical context to which a metaphysical view is held to apply.

Going a step further, we hold that the example of visual perception is a clear case where process externalism applies without reasonable doubt. When this claim is accepted, Kim's critique of mental causation is not only questionable, but arguably false. Mental and neural properties do not refer to the same entity and mental properties cannot be second-order properties of a first-order neural realizer. To repeat, a single counterexample suffices against Kim's assumption that mental and neural properties apply to the same entity, and it seems extremely unlikely that anyone could build a convincing case that such counterexamples are impossible.

A second point that could be raised against our defense of mental causation is that the entity to which mental properties are to be ascribed is not truly mental but actually physical. In this case, the argument goes, mental causation is still threatened by causal exclusion, but now by a conjunction of neural, bodily and environmental processes. However, we want to stress that the component brain, bodily and environmental processes together constitute a single *psychological* first-order entity with its own characteristics. It is to such an extended and micro-based psychological entity that mental properties apply.

A different approach here would be to claim that all new macrocausal powers, such as those occurring in chemistry, biology, but also psychology, are still forms of physical causation (e.g. Kim, 1998, p. 114). Consequently, casting mental causation as another form of macrocausation would not set it apart from mere physical causation. We agree that this is a possible reading of mental causation as here envisioned, but we do not see this possibility as a problem. First, the same would apply to other forms of macrocausation such as occurring in chemistry and biology. When all of this is 'physical', the label loses a lot of its force. Second, psychological phenomena would still provide a genuine addition to the causal structure of the world, they inaugurate new causal powers compared to those of their constituents. Even when these causal powers are considered 'physical' in such a wide and general sense, they would still be specifically mental in being the result of psychological processes rather than anything

else. Third, turning this possible reductive reading of all the special sciences into a mandatory one is work that still needs to be done, and remains rather controversial.

We started out with Kim's gloomy future in which physicalism inevitably leads to what he calls "the irreality of the mental." We hope to have shown that this image does not reflect a deep metaphysical truth, but is based on a premature conclusion that disregards—and is held victim to—discussions and findings relating to the embedded nature of the mind.

Acknowledgements We are indebted for helpful suggestions to Filip Buekens, Dingmar van Eck, Carl Gillett, Susan Hurley, Huib Looren de Jong, Wim de Muijnck, Cory Wright, Marc Slors and other members of the Tilburg/Nijmegen Rationality and Cognition Program and the research group of Theoretical Philosophy in Groningen. Work on this paper was partially supported by the NWO-vidi grant 016.038.301.

References

- Adams, F., & Aizawa, K. (forthcoming). Defending the bounds in cognition. In R. Menary (Ed.), *The extended mind*. Ashgate. <Retrieved from <http://personal.centenary.edu/~kaizawa/cv.htm> on June 7, 2005>
- Baker, L. R. (1993). Metaphysics and mental causation. In J. Heil, & A. Mele (Eds.), *Mental causation*. Oxford: Clarendon Press.
- Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. N. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, 20, 723–742.
- Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72, 173–215.
- Beer, R. D. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Science*, 4(3), 91–99.
- Bennett, M., & Hacker, P. (2003). *Philosophical foundations of neuroscience*. Malden, MA: Blackwell.
- Bickle, J. (1998). *Psychoneural reduction: The new wave*. Cambridge, MA: MIT Press.
- Bickle, J. (2003). *Philosophy and neuroscience: A ruthlessly reductive approach*. Dordrecht: Kluwer.
- Block, N. (2003). Do causal powers drain away? *Philosophy and Phenomenological Research*, 67, 133–150.
- Brooks, R. A. (1999). *Cambrian intelligence*. Cambridge, MA: MIT Press.
- Burge, T. (1993). Mind-body causation and explanatory practice. In J. Heil, & A. Mele (Eds.), *Mental causation*. Oxford: Clarendon Press.
- Chemero, A., & Heyser, C. (2005). Object exploration and a problem with reductionism. *Synthese*, 147, 403–423.
- Churchland, P. M., & Churchland, P. S. (2001). McCauley's demand for a co-level competitor. In W. Bechtel, P. Mandik, J. Mundale, & R. S. Stufflebeam (Eds.), *Philosophy and the neurosciences: A reader* (pp. 457–465). Malden, MA: Blackwell.
- Churchland, P. S. (2002). *Brain-wise: Studies in neurophilosophy*. Cambridge, MA: MIT Press.
- Churchland, P. S., Ramachandran, V., & Sejnowski, T. (1994). A critique of pure vision. In C. Koch, & J. Davis (Eds.), *Large-scale neuronal theories of the brain* (pp. 23–60). Cambridge, MA: MIT Press.
- Clark, A. (1997). *Being there: Putting brain, body and world together again*. Cambridge, MA: MIT Press.
- Clark, A. (forthcoming). Memento's revenge: The extended mind. Extended, In R. Menary (Ed.), *The extended mind*. Ashgate.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58, 7–19.
- De Muijnck, W. (2003). *Dependencies, connections, and other relations: A theory of mental causation*. Dordrecht: Kluwer.
- Dretske, F. I. (1988). *Explaining behavior: Reasons in a world of causes*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1991). Making mind matter more. In *A theory of content and other essays* (pp. 137–159). Cambridge, MA: MIT Press.
- Gozzano, S. (2003). Second order properties: Why Kim's reduction does not work. *Logic and Philosophy of Science*, 1, <<http://www.units.it/~episteme/>>.
- Grush, R. (2003). In defense of some 'Cartesian' assumptions concerning the brain and its operation. *Biology and Philosophy*, 18, 53–93.

- Haugeland, J. (1998). Mind embodied and embedded. In J. Haugeland (Eds.), *Having thought: Essays in the metaphysics of mind* (pp. 207–237). Cambridge, MA: Harvard University Press.
- Hurley, S. L. (1998a). *Consciousness in action*. Harvard: Harvard University Press.
- Hurley, S. L. (1998b). Vehicles, contents, conceptual structure, and externalism. *Analysis*, 58, 1–6.
- Hurley, S. L. (2001). Perception and action: alternative views. *Synthese*, 129, 3–40.
- Hurley, S. L. (forthcoming). Varieties of externalism. In R. Menary (Ed.), *The extended mind*. Ashgate.
- Keijzer, F. A. (2001). *Representation and behavior*. Cambridge, MA: MIT Press.
- Kim, J. (1996). *Philosophy of mind*. Boulder: Westview Press.
- Kim, J. (1997). Does the problem of mental causation generalize? *Proceedings of the Aristotelean Society*, 97, 281–297.
- Kim, J. (1998). *Mind in a physical world: An essay on the mind-body problem and mental causation*. Cambridge, MA: MIT Press.
- Kim, J. (2005). *Physicalism, or something near enough*. Princeton, NJ: Princeton University Press.
- Loewer, B. (2001). Book review of Jaegwon Kim. *Mind in a Physical World. Journal of Philosophy*, 98, 315–324.
- Looren de Jong, H. (2003). Causal and functional explanations. *Theory and Psychology*, 13, 291–317.
- McCauley, R. N. (2001). Explanatory pluralism and the co-evolution of theories of science. In W. Bechtel, P. Mandik, J. Mundale, & R. S. Stufflebeam (Eds.), *Philosophy and the neurosciences: A reader* (pp. 431–456). Malden, MA: Blackwell.
- McClamrock, R. (1995). *Existential cognition: Computational minds in the world*. Chicago: The University of Chicago Press.
- Metzinger, T. (2003). *Being no one: The self-model theory of subjectivity*. Cambridge, MA: MIT Press.
- Noë, A. (2004). *Action in perception*. Cambridge, MA: MIT Press.
- Noordhof, P. (1999). Micro-based properties and the supervenience argument: A response to Kim. *Proceedings of the Aristotelean Society*, 99, 109–114.
- O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24, 883–917.
- Pecher, D., & Zwaan, R. A. (Eds.). (2005). *The grounding of cognition: The role of perception and action in memory, language and thinking*. Cambridge: Cambridge University Press.
- Pereboom, D. (2002). Robust nonreductive materialism. *Journal of Philosophy*, 99, 499–531.
- Pfeifer, R., & Scheier, C. (2001). *Understanding intelligence*. Cambridge, MA: MIT Press.
- Polger, T. W. (2004). *Natural minds*. Cambridge, MA: MIT Press.
- Rensink, R. A. (2002). Change detection. *Annual Review of Psychology*, 53, 245–277.
- Ross, D., & Spurrett, D. (2004). What to say to a sceptical metaphysician: A defense manual for cognitive and behavioral scientists. *Behavioral and Brain Sciences*, 27, 603–647.
- Rowlands, M. (2003). *Externalism: Putting mind and world back together again*. Chesham: Acumen.
- Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, 101, 389–428.
- Schouten, M. K. D., & Looren de Jong, H. (2004). Could the neural ABC explain the mind? *Behavioral and Brain Sciences*, 27, 311–312.
- Simons, D. J. (2000). Current approaches to change blindness. *Visual Cognition*, 1–3, 1–15.
- Simons, D. J., & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences*, 1, 261–267.
- Simons, D. J., & Levin, D. T. (1998). Failure to detect changes to people in a real-world interaction. *Psychonomic Bulletin and Review*, 5, 644–649.
- Thelen, E., & Smith, L. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press.
- Van Gelder, T. (1995). What might cognition be if not computation? *Journal of Philosophy*, 91, 345–381.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin and Review*, 9, 625–636.
- Wilson, R. A. (1995). *Cartesian psychology and physical minds: Individualism and the sciences of the mind*. Cambridge: Cambridge University Press.