To appear in J. Kiverstein and M. Wheeler (eds.), Heidegger and Cognitive Science*, Palgrave-Macmillan*

Naturalizing Dasein and other (Alleged) Heresies

Michael Wheeler

1. This Time it's Personal

To my mind, being wrong is nowhere near as disheartening as being boring, so I am encouraged by the fact that, in the four chapters immediately preceding this one, four thinkers for whom I have nothing but the utmost intellectual respect have found my ongoing project to articulate the philosophical groundwork for a genuinely Heideggerian cognitive science interesting enough that they have taken the trouble to explain precisely why it is flawed. Just how deep the supposed flaws go depends on which set of criticisms one chooses to read. For Ratcliffe and Rehberg they go very deep indeed, since, for these thinkers, there is a sense in which the very idea of a Heideggerian cognitive science borders on the incoherent. Dreyfus and Rietveld, on the other hand, seem to agree with me that something worth calling a Heideggerian cognitive science is certainly possible; it's just that my version of it is seriously defective.

Although being interesting and wrong is preferable to being boring and right, one aspires to be interesting and right. In what follows, then, I shall endeavour to resist what I take to be the most powerful (although admittedly not all) of the criticisms tabled in each of the chapters just mentioned. In setting up my contribution to the present volume in this way, I don't for one moment mean to suggest that criticizing my position is all, or indeed the most important thing, that these papers do. Far from it: the positive contributions they make in bringing the relations between Heideggerian philosophy and cognitive science into better view – the multiply elusive 'and' in the identifier 'Heidegger and cognitive science', as Rehberg (this volume) would have it – are a great deal more significant than their problematizations of my own view. Still, I am self-interested enough to spend my time here defending myself in the line of fire, which is why, in what follows, I shall be arguing that Ratcliffe, Rehberg, Dreyfus and Rietveld (in roughly that order) have further work to do before my letter of surrender arrives in their inboxes. That said, the goal of this chapter may also be expressed in a way that avoids the egocentric emphasis placed on it so far: that goal is to defend both the very idea of a Heideggerian cognitive science and a certain vision of what such a cognitive science will look like, at least in part. So, adopting something of a pantomime register, I shall suggest that the right response to Ratcliffe's (this volume) strident claim that there can be no cognitive science of Dasein (where 'Dasein' is Heidegger's term for the distinctive kind of entity that human beings as such are) is to shout back (with gusto) 'oh yes there can!'.

2. Keeping Dasein out of the Lab

Let's begin with the arguments of Ratcliffe and Rehberg. It is important to record at the outset that neither of these thinkers hold cognitive science to be an intellectually worthless endeavour that is incapable of yielding important insights about its subject matter. Thus Rehberg (this volume, p.??) stresses that "[w]hat is at stake… is not scientific practice, scientific achievements or the positive role the sciences play in modern life", while Ratcliffe (this volume, p.??) emphasizes that his intention "is not to

dismiss cognitive science altogether but to suggest that there is a principled limit to its potential aspirations". In their different ways, however, Ratcliffe and Rehberg each argue that there exists some sort of *in-principle* barrier to the proposal that cognitive science might be developed in a genuinely Heideggerian direction. Hence Ratcliffe's aim, as he himself describes it, is "to raise some philosophical concerns about the very idea of a 'Heideggerian cognitive science'" (this volume, p.??). From this perspective, then, there is no need to examine the specific explanations that any so-called Heideggerian cognitive science might have offered – or at least if one does examine them one shouldn't take their purported Heideggerian character seriously – because, in truth, the whole nascent paradigm is a kind of fiction.

The first volley in Ratcliffe's broadside is a general argument to the effect that Heidegger's philosophy is incompatible with the naturalism that, according to me (see e.g. Wheeler, 2005, pp.4-7), ought to, and standardly does, accompany cognitive science. Indeed, it has always seemed to me that a healthy respect for the rich and diverse research programme that we call cognitive science requires a fundamental commitment to a thoroughgoing naturalism regarding human psychological phenomena. One can put this point another way: any philosophy of mind and cognition that rides shotgun with cognitive science must be naturalistic in form. So what is it for philosophy to be naturalistic in form? The guiding thought of naturalism is that philosophy should be *continuous with* empirical science. The question of just what this continuity might amount to will become an important issue for us soon. At present let's simply note that the naturalist about some phenomenon X (e.g. about mind, cognition, sense-making or being-in-the-world) holds that the science related to X places constraints on our philosophical theorizing about X. Now, a genuinely Heideggerian cognitive science (as opposed to a cognitive science that occasionally borrows Heideggerian insights) will be a cognitive science that, to some significant degree, systematically integrates Heidegger's philosophical framework with the foundational features of the cognitive-scientific approach to mind, intelligence, thought and action. So if, as I have suggested, taking cognitive science seriously requires a commitment to a naturalism about psychological phenomena, then the prospects for a Heideggerian cognitive science rest, in part, on whether or not Heideggerian philosophy is, or can be made, compatible with that naturalism.

It is at this point that Ratcliffe bares his teeth. According to Heideggerian philosophy, as interpreted by Ratcliffe, the distinctive manner in which empirical science, and so cognitive science in particular, reveals entities as the targets and the outcomes of its investigations, tacitly presupposes a sense of *belonging to the world* on the part of Dasein. This sense of belonging to the world is, in effect, a dimension of the distinctive manner in which Dasein is essentially *in the world*, an 'in-ness' that Heidegger christens *dwelling* (see e.g. Heidegger, 1927, pp.79-80). To dwell in a house is not merely to be inside it spatially in a physical sense. Rather, it is to *belong* there, to have a *familiar place* there. This belonging or familiarity is sometimes illuminated by Heidegger by way of the thought that entities ordinarily make sense (are intelligible) to us within culturally and historically determined contexts of practical activity that, so to speak, arrive *with* us. For example, my laptop currently makes sense to me in relation to a skilled activity of text-editing; that text-editing is involved in writing a document; that document-writing is involved in meeting a professional deadline; and that meeting of a professional deadline is involved in *my* project of being a good academic. In a fundamental sense (although see below for a qualification), it is these Dasein-relative structures of significance within which entities are found that, in Heidegger's analysis, interconnect, combine and interweave to make up the meaningful structure of a world. And what this indicates is that, in practical encounters with entities – encounters in

which, for Heidegger, entities are encountered principally as *ready-to-hand* (i.e., smoothly available for skilled activity; see below) – the world is something with which each of us (as Dasein) is always already familiar, and to which each of us (as Dasein) has a sense of belonging (see e.g. Heidegger, 1927, p.119).

As Ratcliffe understands things, Heidegger's account of our belonging to the world generates a roughly transcendental case against naturalism, and so neatly places an in-principle barrier in the path of any proposal for a Heideggerian cognitive science. To see how this is supposed to work, we need to begin by reminding ourselves of Heidegger's famous claim that science reveals entities in a mode of being (intelligibility) that he calls *presence-at-hand*. When revealed as present-at-hand, entities are encountered as removed from the familiar settings of everyday practical activity and thereby emerge as the kind of context-independent objects that populate science, that is, as the bearers of certain cross-contextual determinate or measurable properties (size in metres, weight in kilos, position in objective space etc.). With presence-at-hand brought into view, there are in truth two different ways in which a transcendental bulwark against naturalism might be erected, based on two different notions of worldliness that appear in *Being and Time*. (Although Ratcliffe doesn't quite present things in this way, I do not think he would have any serious objections to my restaging, since, as I shall indicate, both of the resulting arguments appear in his text.) If one thinks of Dasein's world as constituted by some global network of the kind of local contexts of practical activity highlighted above, then one might very well argue that the entities of science are "presented to us *as* stripped to varying degrees of the significance that they previously had, as having lost something" (Ratcliffe, this volume, p.??). On this view, to experience an entity as present-at-hand is to encounter it precisely as having-been-removed-from-the-world, which is of course a possible style of encounter only if the world itself (the network of Dasein-relative contexts of practical activity with which we are familiar) is presupposed (cf. Wheeler, 2005, p.165). Alternatively (and this is the aforementioned qualification to the Heideggerian notion of 'world'), if one thinks of Dasein's world not as a network of interconnected contexts of practical activity, but as the very structure of intelligibility itself (what Heidegger sometimes calls the *worldhood* of the world; see e.g. Heidegger, 1927, p.119), and if one thinks of scientific practice as itself a kind of sense-making, then one might very well argue that "both present-at-hand and ready-to-hand entities presuppose a world, *within which* it is possible to encounter entities in these ways" (Ratcliffe, this volume, p.??, my emphasis). And then it appears that the sense of familiarity and belonging that characterizes our distinctive world-embeddedness will inevitably accompany scientific sense-making. Either way, then, it seems eminently arguable that Dasein's sense of belonging to the world is a transcendental condition of the distinctive mode of sense-making that is characteristic of cognitive science, and so, one might think, cannot be brought within the explanatory reach of that science.

Like Ratcliffe, Rehberg argues that, from a Heideggerian perspective, there is something about the fundamental philosophical profile of cognitive science, construed as a local dimension of modern scientific thinking in general, that renders it incapable in principle of reaching its explanatory goal. Once again the Heideggerian claim, that it is of the essence of modern science to reveal entities as present-at-hand, is to the fore, but now that claim takes on a form more readily associated with the later Heidegger (e.g. Heidegger, 1954) than the Heidegger of *Being and Time*. Articulated in this later modulation, the objectification of entities on which modern science depends – realized as a process of mathematization that renders entities apt for measuring and technical manipulation – is not only itself a partial understanding of entities (one mode of sense-making or being among others), it also obscures first, the entity as a site of multiple alternative dimensions of sense-making (what Rehberg

calls the *self-differing* nature of entities), and second, the essential revealing-concealing dynamic of that sense-making. In other words, modern science has the property of obscuring the fact that any way of making sense of entities (including its own) involves, as its concomitant flip-side, a necessary concealing of the plenitude of other (e.g. cultural, religious) ways in which those entities may have become intelligible. It is through this doubly obscuring character, based on the reduction of entities to objects and of intelligibility to the measurable and the manipulable, that scientific thinking comes to present itself as the one dominant account of reality, excluding all others.

If Heidegger is right, then the philosophical irony, of course, is that science, as a mode of revealing, depends ultimately on the revealing-concealing dynamic of being, so science obscures the fundamental structure of its own functioning. But in the present context it is perhaps more telling to stress the related point that the scientific reduction of entities to mere objects ultimately obscures the fundamental self-differing character of entities. Thus "the more vigorously a science pursues its object, the more does the being, which it thus attempts to capture as object, withdraw itself" (Rehberg, this volume, p.??). The explanatory 'object' of a Heideggerian cognitive science would, of course, be Dasein as being-in-the-world, which generates the following, local application of the more general claim:

> the more comprehensively cognitive science tries to secure Dasein as its object, the more surely will Dasein escape objectification. But if 'cognitive science' were capable of… giving up its essence – it would be able to approach Dasein without reducing it to the status of an object, but then it would not be a science, and it would not focus on human cognition – but perhaps on its ways of being. (Rehberg, this volume, p.??)

According to Rehberg, then, a genuinely Heideggerian perspective is committed to the principle that Dasein (human being-in-the-world) may be understood only if we refrain from reducing it to a mere object. Since Dasein is the explanatory target of a Heideggerian cognitive science, for such a cognitive science to succeed, it would need to adopt a non-objectifying mode of sense-making. But the objectifying mode of sense-making is an essential characteristic of modern science and thus of cognitive science. So a cognitive 'science' that adopted a non-objectifying mode of sense-making, and thereby managed to approach its explanatory target in a potentially fruitful way, would simultaneously cease to be a science. Doubt is thus cast on the very possibility of a genuinely Heideggerian cognitive science.

(The observant reader will have noticed that there is a second aspect to Rehberg's argument, as summarized in the quotation just above, namely that a cognitive 'science' that adopted a non-objectifying mode of sense-making would, in that process, shift its focus so as to concentrate not on human cognition but on Dasein's situated ways of being. Although I do not have the space in the present treatment to pursue this thought in detail, it seems to me that, to the extent that there is an issue to be addressed here, it turns on the already-highlighted point about what is essential to cognitive science as a science. After all, the Heideggerian cognitive scientist is going to want to eschew narrow or Cartesian delineations of the domain of the cognitive, and to include Dasein's situated ways of being within that domain, thereby agreeing to the shift in focus. So there is only a problem here if the Heideggerian cognitive scientist is somehow prevented from making that transition. One might argue that this would be the case, if the transition in question is impossible without Heideggerian cognitive

science surrendering its commitment to objectification (mathematization, measurability) and thus its status as a science. In a more indirect, but correspondingly richer way, Rehberg herself makes something like this point (see e.g. Rehberg, this volume, pp.??).)

That, then, is the in-principle case against a Heideggerian cognitive science that I shall consider here. If Ratcliffe and Rehberg are right, the very idea of a Heideggerian cognitive science is fundamentally flawed. But are they right?

3. Heideggerian Naturalism

Ratcliffe and Rehberg present arguments that differ in their details. Nevertheless, those arguments have something in common, namely the thought that, for the scrupulous Heideggerian, there is something about cognitive science (whether it be a presupposed sense of belonging to the world or an objectifying mode of sense-making) that renders that discipline incapable, in principle, of illuminating human psychological existence in the manner to which it aspires. This shared thought may help to explain why, as we are about to see, both lines of critique are in fact blocked by the very same structural feature of the relationship that obtains between, on the one hand, the philosophical illumination of human thought and activity and, on the other, the cognitive science of that same thought and activity. This feature is one that Heidegger himself arguably explicated as clearly as anyone has (although, in truth, that amounts to nowhere near as clearly as one would like). Put bluntly, and in a way that makes the point seem largely banal, the key observation is that although, when it comes to the study of human being, the deliverances of philosophy and the deliverances of empirical cognitive science sometimes constrain each other in intelligible ways, the two disciplines are concerned centrally with generating different *kinds* of understanding. As Gallagher and Zahavi (2008, p.7) put it, with a focus on phenomenological philosophy that anticipates where we are going here, the phenomenologist and the psychologist are "taking different approaches, asking different questions, and looking for different kinds of answers".

I shall attempt to do justice to these differences by adopting a version of McDowell's (1994) distinction between *constitutive* understanding and *enabling* understanding. Constitutive understanding is the characteristic target of philosophy, although not only of philosophy. It concerns the identification, articulation and clarification of the conditions that determine what it is for a phenomenon to be the phenomenon that it is (e.g. what it is for a certain kind of creature to competently inhabit its world). Enabling understanding is the characteristic target of empirical science, although not only of empirical science. It reveals the causal elements, along with the organization of and the systematic causal interactions between those elements, that together make it intelligible to us how a phenomenon of a certain kind could be realized or generated in a world like ours (e.g. how some creature-specific mode of competent world-inhabiting is causally enabled in a purely physical universe). I shall argue that once we bring these two kinds of understanding, and the relationship between them, into proper view, the tension between Heideggerianism and cognitive science evaporates. Moreover, it evaporates in a manner that not only renders Ratcliffe's point about our sense of belonging to the world consistent with a perfectly respectable psychological naturalism, but also allows a Heideggerian cognitive science to absorb Rehberg's point about the objectifying effects of science. Unusually, then, this is a case in which it is possible to have one's cake and to eat it.

(It is worth noting that McDowell tends to think of the constitutive understanding of human mentality and experience as being pitched at what he calls the *personal* level, and of the enabling understanding offered by empirical cognitive science as being pitched at what he calls the *subpersonal* level (McDowell, 1994). I avoid this additional terminology here, because, without the inclusion of a long and detailed analysis, the notion of the personal level would seem to imply the existence of some self-sufficient ontological entity – i.e., the person – that threatens to clash with Heidegger's process-based notion of Dasein. For similar reasons, I am happy to concede Rehberg's point that my previous talk of 'Dasein' as another term for the human agent is apt to be misleading (Wheeler, 2005; Rehberg, this volume). For what it's worth, I think that an analysis that relieved the tensions here could be given, but in the present context it is more appropriate and less onerous simply to avoid the troublesome terminology.)

Heideggerian philosophy may be depicted as seeking a particular sort of constitutive understanding. Articulated in a way that deliberately ignores certain well-rehearsed, politically charged issues of chauvinistic prejudice that are sometimes thought to attach to Heidegger's analysis (see e.g. Wolin, 1993), the understanding in question concerns an account of the conditions that determine what it means to live a human life. Heidegger's philosophical strategy for delivering such an account of human existence is, famously, *phenomenological* in character. As developed in *Being and Time* (and arguably throughout Heidegger's writings; see e.g. Wheeler forthcoming), phenomenology may be depicted as a theoretical (or perhaps a meta-theoretical) philosophical enterprise that, through an attentive and sensitive examination of ordinary human experience, aims to reveal the transcendental yet historical conditions which give that experience its form. Because these target structures are transcendentally presupposed by ordinary experience, they must in some sense be present with that experience, but they are not simply available to be read off from its surface, hence the need for disciplined and careful phenomenological analysis to reveal them. The historicality exhibited by the transcendental here is a consequence of what Heidegger takes to be the hermeneutic character of understanding in general, and thus of phenomenological understanding in particular. As an interpretative activity, phenomenological analysis is inevitably guided by certain historically embedded ways of thinking that the phenomenologist brings to the task, meaning that its results remain ceaselessly open to revision, enhancement and replacement.

What results from Heidegger's phenomenological investigation of human existence is, of course, his ground-breaking account of being-in-the-world. Here is not the place to attempt to lay out the systematic structure or content of that account (see Wheeler, forthcoming, for my own introduction), although later in this chapter some of its salient features will be discussed. What matters at present is the relationship that obtains between that constitutive understanding of human existence (whatever its details) and cognitive science. As already noted, cognitive science is paradigmatically in the business of providing an enabling understanding of psychological phenomena. More specifically, the goal of cognitive science is to map out the causal elements whose organization, operation and interaction enable psychological phenomena. Viewed through this philosophical lens, the aim of a Heideggerian cognitive science would be to develop an enabling understanding of (i.e., an account of the causal states and processes that underpin) the phenomena that constitute Dasein's being-in-the-world. Again, for the moment, the details of such an account don't matter, although later in this chapter we will meet some examples of cognitive-scientific models that plausibly contribute to it. What concerns us right now is the relationship that obtains between

Heidegger's phenomenological constitutive account (whatever its details) and the corresponding cognitive science (whatever its details). What, then, is that relationship?

There is no good reason to think that constitutive understanding and enabling understanding must be reductively related. In particular, there is no good reason to think that constitutive accounts *must* be reducible to enabling accounts. That would be tantamount to a kind of runaway *scientism* (more on which in a moment), unpacked as the thought that science is the measure of all things. But, equally, there is no good reason to think that constitutive understanding and enabling understanding *must* be wholly independent of each other. Rather, our two kinds of understanding (and thus philosophy and cognitive science) will standardly engage in a process of mutual constraint and influence that McDowell (1994, p.197) tags with the enticing phrase "a perfectly intelligible interplay". Although McDowell himself says disappointingly little about the nature of this interplay, its general shape seems clear enough.

On the one hand, there is a sense in which philosophy, as a source of constitutive understanding, will isolate and clarify phenomena for which the corresponding cognitive science will then try to identify the underlying causal mechanisms. According to this model of constitutive-to-enabling influence, a phenomenological analysis that maps out the essential features and the conditions of possibility of human experience will provide a rich and systematic account of what the science of consciousness needs to explain. To borrow an example from Gallagher and Zahavi (2008, p.10), if disciplined phenomenological analysis suggests an experiential profile according to which perception is always perspectivally incomplete (i.e., we never see all of an object at once), even though objects are presented to us in perception as having aspects that, right now, we cannot see, then any science of consciousness must respect and account for that profile.

On the other hand, the causal profiles discovered by cognitive science may sometimes lead us to revise our conception of what the phenomena under investigation are. McDowell's (1994) own example of such enabling-to-constitutive influence is (something like) the following. (I say 'something like' because I have assembled the example I am about to give out of more than one passage in McDowell's discussion.) Imagine that there is some creature that, as far as our pre-scientific observations of its behaviour go, seems to enjoy a rich kind of perceptual contact with the world. Scientists then open up the head of that creature only to discover that there is nothing more complex driving its behaviour than a lump of homogenous jelly. Such a discovery would arguably be followed by a downgrading, within our constitutive understanding, of the creature's world-inhabiting competence. A less dramatic (but comfortingly actual) example of enabling-to-constitutive influence (cited by Cappuccio and Wheeler, 2010) concerns recent experiments on mirror neurons. These experiments have influenced the phenomenological understanding of motor intentionality and intersubjectivity by suggesting that an agent's own motor competences enrich and actively modulate her perception, recognition and categorization of intentional actions, endowing her with a prereflective understanding of the intentions of other agents (Rizzolatti and Sinigaglia, 2007; see De Preester, 2008, for an analysis in which a consideration of mirror neuron research is used explicitly to drive the phenomenological-level claim that Merleau-Ponty's account of self-other understanding as world-mediated presupposes a Husserlian notion of pairing or bodily similarity). Enabling-to-constitutive influence in the vicinity of phenomenology has even been incorporated explicitly into a proposed research methodology. In Gallagher's (2003) *front-loaded phenomenology*, the results of phenomenological analysis are tested experimentally, and the experimental outcomes are then used to confirm, refute, revise or extend that phenomenological analysis. (Gallagher's specific example concerns the empirical

testing of a phenomenologically established distinction between one's sense of being the agent of an action and one's sense of ownership of one's body; see (Gallagher and Zahavi, 2008, p.39).)

About now, some readers of this chapter will want to complain that none of my examples of the phenomenological engagement with cognitive science concern Heidegger's own hermeneutic phenomenological framework, a framework which, as Ratcliffe and Rehberg have shown us, might well engender hostile diplomatic relations with cognitive science. Time, then, to address this shortfall.

Card carrying Heideggerians and their intellectual brethren who have reflected explicitly on the possibility of a *productive* relationship between phenomenology and cognitive science have sometimes argued for a one-way constraint from the former to the latter, on the sort of grounds already mentioned, namely that since "phenomenology tells us about the phenomena of human experience, its results ought to be relevant to the human sciences" (Kelly's report of a view expressed by Dreyfus; Kelly, 2000, p.161). I suggest, however, that Heidegger himself, in relation to the scientific research on human experience and behaviour that was available in his time (branches of anthropology, psychology and biology), not only made ample conceptual room for something very like the two-way intelligible interplay advocated by McDowell, he also had a more developed account than McDowell of how that interplay is supposed to work. The interpretation of Heidegger that I am about to reiterate (it appears in a somewhat less developed form in Wheeler, 2005) is, I freely admit, controversial in places, but I cannot help that. It seems to me to be the only way to make sense of the conjunction of the following two passages from *Being and Time*, each of which, in my view, isolates one of the different directions of influence (constitutive-to-enabling, enabling-to-constitutive) that the McDowellian interplay involves.

> [The] ontological foundations [of anthropology, psychology, and biology] can never be disclosed by subsequent hypotheses derived from empirical material... they are always 'there' already, even when that empirical material simply gets collected. (Heidegger, 1927, p.75)

> [T]he positive sciences neither 'can' nor should wait for the ontological labours of philosophy to be done… the further course of research will not take the form of an 'advance' but will be accomplished by recapitulating what has already been ontically discovered, and by purifying it in a way which is ontologically more transparent. (Heidegger, 1927, p.76)

Let's consider these passages in turn. To understand the first we need to recognize that, for Heidegger, ontological analysis is the making explicit of the essential structure and the conditions of possibility of the phenomena of human experience. In other words, it is the kind of constitutive analysis carried out by phenomenology. Moreover, according to Heidegger, any empirical science (and that includes anthropology, psychology, biology or cognitive science) will be structured in such a way that, in order to deliver its distinctive species of enabling understanding, it must assume certain basic concepts and principles – the ontological foundations – that determine the constitutive character of its target phenomena. In much cognitive science, for example, some notion of internal representation is assumed that shapes the subsequent empirical study of mind. These ontological foundations, Heidegger tells us, cannot *simply* amount to further empirical hypotheses. That is, they cannot be delivered *in a fully articulated form* as a *straightforward* output of a distinctively scientific

methodology. This makes prima facie sense, since those foundations are presupposed by the empirical research in question. Crucially, however, it does not follow, from the fact that the ontological foundations of a science are not on all fours with the ordinary empirical hypotheses of that science, that those foundations are somehow wholly independent of the ongoing empirical research. And that's where the second passage from *Being and Time* comes in.

Here is something that Heidegger doesn't quite say in that second passage, but which I take it is strongly suggested by his talk of empirical science not needing to wait for philosophy and of phenomenological research (which is how I read the phrase "the further course of research" – see below) as delivering an ontologically purer form of what has already been ontically (i.e. scientifically) discovered. As we have seen, for Heidegger, any particular example of empirical science can take place only on the basis of certain ontological presuppositions. If the presuppositions in play are fruitful generators of predictively and explanatorily powerful empirical models, then the ordinary dynamic of scientific research will maintain their presence at the heart of that science, as something akin to the defining concepts and principles of a Kuhnian paradigm, or the hard core of a Lakatosian research programme (Kuhn, 1962; Lakatos, 1977). If, by contrast, the presuppositions in play become identified as the source of stalled empirical models that consistently fail to account for new or historically recalcitrant data, then the science itself will tend to revise those presuppositions, perhaps in a revolutionary, paradigm-shifting transition. Kuhn and Lakatos said that too. If all this is right, one might wonder what there is left for philosophy to do. But nothing about the proposed picture suggests that the ontological foundations of a science cannot be strengthened or criticized by philosophical (e.g. phenomenological) argument. Moreover, according to Heidegger (and now we're back with the first of the target quotations), such foundations will tend to remain buried or impure without the benefit of philosophical (phenomenological) amplification and clarification.

The overall intellectual dynamic just sketched, which, as I have argued, is implied by Heidegger's remarks, requires that any particular set of science-shaping ontological foundations (any particular presupposed constitutive understanding) must, in principle, be sensitive to, and modifiable in the wake of, the success or failure of the empirical models which those foundations support. In other words, for Heidegger, good empirical science (good enabling understanding) depends, in part, on good ontological foundations (good constitutive understanding), and while the ontological foundations of a science may be articulated, clarified and criticized by philosophy (by phenomenology), they may also be revised in the light of ongoing empirical scientific research. What this tells us is that the phenomenology-involving version of the McDowellian intelligible interplay, the structure for which we have been searching, is indeed fully present in Heidegger's framework. Moreover, if we squint just hard enough, Heidegger may be seen as recommending a broadly Kuhnian or perhaps Lakatosian model for how that interplay functions.

One might, I suppose, object to my analysis here by challenging my reading of 'research' in the target quotation as indicating phenomenological research, but if one reads it as indicating scientific research instead, it is hard to make sense of the tasks of purification and increasing transparency that, Heidegger argues, will achieve the 'advance' in question. In Heidegger's thinking those tasks certainly seem to be the responsibility of philosophy, and in particular of ontological analysis. A more substantive objection would claim that I have ignored a critical aspect of Heidegger's framework, by considering regional ontology (focussed on the sets of ontological foundations that determine how sense-making works in local contexts, for

example within particular sciences) in isolation from its dependence on fundamental ontology (focussed on what it is that unites and makes possible all our varied and diverse modes of sense-making, including our regional ontologies). Heidegger is laudably clear that regional ontology depends on fundamental ontology. As he puts it, "the question of Being aims… at ascertaining the a priori conditions not only for the possibility of the sciences which examine beings as beings of such and such a type, and, in doing so, already operate with an understanding of Being, but also for the possibility of those ontologies themselves which are prior to the ontical sciences and which provide their foundations" (Heidegger, 1927, p.31). But if regional ontology depends ultimately on fundamental ontology, then for all I've said there remains an empirically inviolable foundation to distinctively human existence (the human capacity for sense-making) that lies beyond the reach of cognitive science. I am unmoved by this supposed worry, in part for general reasons concerning the commitments of naturalism that I will discuss in a moment, but also because Heidegger is adamant that the basic condition for intelligibility that fundamental ontology will articulate does not exist as an ethereal metaphysical something that lies beyond entities in the realm of what Sheehan (2001) calls "Big Being". Rather, in a catchy little Heideggerian phrase, being is always the being of some entity. There is much to be said about what this claim means, but I see no way of making sense of it while simultaneously insulating the results of fundamental ontology from the results of regional ontology. And given that, on Heidegger's account, regional ontology is subject to empirical impact (although not, of course, to scientific reduction), fundamental ontology must be too.

So, according to me, the Heideggerian framework, when interpreted correctly, embraces a neo-McDowellian intelligible interplay between constitutive understanding or, more specifically, phenomenology, and enabling understanding or, more specifically, cognitive science. Say I'm right. What are the consequences for the arguments that Ratcliffe and Rehberg present against a Heideggerian cognitive science?

To hold Ratcliffe at bay, the emerging Heideggerian picture must be rendered compatible with naturalism. As noted earlier, the guiding thought of naturalism is that philosophy should be *continuous with* empirical science. To evaluate whether or not we have the desired compatibility, then, we need an answer to the following question: how is the notion of continuity that is at the heart of naturalism to be understood? There is no univocal answer to this question. Rather, the stripe of one's naturalism is determined by the particular account of continuity that one gives. For example, a certain kind of hard-headed naturalist about psychological phenomena might understand continuity in terms of the reduction of pre-scientifically identified psychological phenomena to scientifically identified states and processes. If, along with thinkers like McDowell (1996), one thinks of pre-scientifically identified psychological phenomena as occupying a logical space of reasons (of justification), and of science as offering explanations that come from outside that space (maybe, from a realm of law, although see McDowell, 2009), then the reductionist claim that one can give an exhaustive account of the former in terms of the latter generates what McDowell (1996, p.67) calls *bald naturalism*. However, this strong reading of continuity is not mandatory. For example, we might read continuity with natural science in what is perhaps the weakest possible way, that is, as mere *consistency with* empirical science (Wheeler, 2005). In this case we avoid any blanket demand for scientific-reductionist explanations of psychological phenomena, while allowing for such explanations in specific cases. (For a pretty much equivalent conception of naturalism, see Elton's (2003) interpretation of Dennett.) Let's call this position *minimal naturalism*. (There is an interpretative question that I don't examine here, as to how the minimal naturalism that I favour, and that I take to be compatible with

McDowell's account of the relationship between constitutive and enabling explanation, relates to the brand of naturalism to which McDowell himself subscribes, a brand that turns on the claim that although human sensibility exists within nature, a complete constitutive account of that sensibility cannot be given using concepts from outside the logical space of reasons (McDowell, 2006); for critical discussion of McDowell's naturalism, see e.g. (Haddock, 2009).)

Despite its deliberately modest character, minimal naturalism still has enough bite to count as a genuine form of naturalism. That's because the consistency condition, as I understand it, involves a constraint running from science to philosophy, to the effect that if and when there is a genuine clash between philosophy and empirical science (in the sense that philosophy demands the presence of some entity, state, or process which is judged to be inconsistent with empirical science), then it's philosophy and not science that must ultimately concede, through withdrawal or revision of its claims. The naturalistic pressure that cognitive science exerts on phenomenology is just a special case of this constraint. (My way of unpacking the consistency condition owes a debt to Huw Price's formulation of what he calls *subject naturalism* as being the view that "[s]cience tells us that we humans are natural creatures, and if the claims and ambitions of philosophy conflict with this view, then philosophy needs to give way" (Price, 2004, p.4). This is not to claim, of course, that Price's subject naturalism is equivalent to my minimal naturalism; it is not.) I include the qualification 'ultimately' in my formulation because not even the most enthusiastic naturalist should expect good philosophy to capitulate to bad science. Strictly speaking, then, the constraint on the table ought to be that if there is a clash between philosophy and some *eminently well-supported* science, then it is philosophy that should concede, through the withdrawal or revision of its claims, although there will often be room for negotiation. This constraint secures, but needs nothing more radical than, the kind of enabling-to-constitutive influence that the Heideggerian version of the McDowellian intelligible interplay requires. But what about the other direction of travel? Is minimal naturalism compatible with the kind of constitutive-to-enabling influence that Heideggerian thinking demands?

The first thing to note is that although minimal naturalism makes cognitive science *relevant to* our phenomenological account of Dasein, it does not usher in cognitive *scientism*. In other words, because it stops a long way short of reductionism, minimal naturalism does not demand that a complete cognitive science of Dasein would be a complete understanding of Dasein, although it would be a complete enabling understanding. I take it that this is 'merely' a more general expression of the thought that all right-thinking philosophers and psychologists have concerning, say, evolutionary-psychological theorizing about the mechanisms that underlie our moral reasoning. One can surely hold that such theorizing tells us much that is important about the way human beings reason about ethical issues, including by suggesting, or by placing limits on, the kinds of factors to which a constitutive account ought to count us as being sensitive, without also holding that evolutionary fitness determines what a human being should – or even, given rational argument or other contextual influences, will – judge to be morally correct. In short, where a constitutive claim could, in principle, clash with science, the consistency constraint will operate, but not all claims are like that. So, with continuity understood in terms of consistency, the fact that there are principled limits on what cognitive science might explain is no longer a threat to naturalism, so it is no longer a threat to a Heideggerian cognitive science.

It might seem that this argument doesn't go quite far enough. After all, Ratcliffe's argument cites a phenomenon, our sense of belonging to the world, that, according to Heideggerian

thinking, is presupposed by cognitive science, in the sense that it is a transcendental condition of the distinctive mode of revealing entities that is characteristic of cognitive science. It is this very transcendality that is supposed to make Dasein resistant to naturalistic understanding. In fact, however, the threat-level is not increased by this reminder. According to the intelligible interplay view that I have attributed to Heidegger, the conditions that are transcendentally presupposed by any particular scientific explanation or model (the structures and elements of regional and fundamental ontology), are not immune to revision, or even perhaps rejection, in the light of the results of the empirical scientific research that they make possible. Indeed, while there might well be philosophical notions of the transcendental that succeed in screening off the conditions of possibility of human experience from scientific influence altogether, it is mysterious how the domesticated and historicized notion of the transcendental that functions in Heidegger's phenomenology could have that effect. After all, according to Heidegger, the conditions of possibility of specific enactments of sense-making are concretely embedded in history (historicality is part of the existential constitution of Dasein), and thus may change as societies change, including, presumably, as a result of scientific change, given that science is itself a sense-making activity embedded in human history. So, in the light of the feedback loops in play, it is perfectly consistent with Heideggerian thinking to hold that the details of our phenomenological understanding of our sense of belonging to the world could be reshaped in the light of cognitive-scientific research on the enabling conditions that causally generate that sense of familiarity. In the next section we shall begin to investigate the cognitive science of belonging to the world.

Before that, however, we need to revisit Rehberg's claim that the objectifying mode of sense-making which, according to Heideggerian thinking, is an essential characteristic of cognitive science, renders cognitive science inherently incapable of illuminating Dasein's being-in-the-world. The minimally naturalistic Heideggerian picture that I have been painting deflects this objection too. To see how, it will be useful to reconstruct Rehberg's claim in the language of the distinction between constitutive understanding and enabling understanding which, as I have argued, is a keystone of Heidegger's own reflections on phenomenology and the human-related sciences. Put in such terms, the Heideggerian viewpoint is that the dominant Western way of inhabiting science is one in which the objectifying explanations of cognitive science are offered as constitutive accounts of human activity, that is, as entailing a general, cognitive-*scientistic* reduction of phenomenological (constitutive) understanding to scientific (enabling) understanding in the vicinity of human psychological life. It is certainly true that, on Heideggerian grounds, we should resist this kind of general objectifying reduction as obscuring the essential character and dynamics of human sense-making. But that is not to resist the idea of a Heideggerian cognitive science. Far from it. It is to resist a certain *manner of inhabiting* cognitive science, a manner that is not mandated by the minimal naturalism that I have argued will accompany any cognitive science with genuinely Heideggerian credentials. It is possible for the objectifying, causal-enabling explanations of human behaviour that cognitive science delivers to inform phenomenological analysis in the fashion required by our naturalistic intelligible interplay without there being any threat that what is thereby established is the kind of objectified levelling of human being against which Heidegger warns us. Moreover, this is all possible without cognitive science surrendering any of its essential characteristics as part of modern science, and thus ceasing to be a science.

The first round of my self-interested defence is now complete, with the very idea of a Heideggerian cognitive science emerging largely unscathed. Round two, in which I shall respond to criticisms by Dreyfus and Rietveld, will take a rather different form. That's because Dreyfus and Rietveld agree with me that a research programme worthy of the name

'Heideggerian cognitive science' is possible. The problems that they wish to lay at my door concern the details of my particular account.

4. Dasein in the Lab: the Heideggerian Cognitive Science of Relevance Sensitivity

What would a Heideggerian cognitive science look like? That, of course, is far too big and complicated a question to address in the few pages I now have available. So here is a substitute: what would a Heideggerian cognitive science *of human sensitivity to relevance* look like? In other words, from a Heideggerian perspective – indeed, from within the kind of minimally naturalistic Heideggerian framework sketched above (or something essentially similar) – what are the causal mechanisms that enable human sensitivity to relevance? Why is this an interesting question? The first thing to note is that even in the sort of dynamically shifting and open-ended scenarios in which we often find ourselves, human beings are extraordinarily proficient at maintaining psychological and behavioural focus on what is contextually relevant in a situation, while ignoring what is contextually irrelevant. For the most part anyway, we think and act in ways that are fluidly and flexibly keyed to context-dependent relevance. Since this impressive capacity is undoubtedly central to human intelligence, cognitive science should be able to arrive at a credible enabling understanding of it – or rather, if minimal naturalism is correct, it should. The fact is, however, that a range of prominent thinkers in and around cognitive science, from Fodor (1983) at the classical end of the spectrum to Dreyfus (e.g. 1990, 1992, this volume) at the Heidegger-influenced other end, have claimed that cognitive science has historically failed to meet this expectation.

The empirical evidence often cited in favour of this judgment is the seemingly recalcitrant nature of two relevance-driven problems – the commonsense knowledge problem and frame problem – that have arguably plagued work in artificial intelligence (see e.g. Dreyfus, this volume). For our purposes, what matters is that these problems plausibly have a similar structure. Ignoring many important details, that structure looks like this. Orthodox cognitive science is built on the claim that intelligence is a matter of building, retrieving, manipulating and transforming suites of representational states and processes. Against this theoretical backdrop, the natural thought when it comes to the issue of relevance is that the intelligent agent should specify and track relevance, by systematically representing the key features of the contexts in which she finds herself. These context-specifying inner representations will in turn determine which first-order inner representations are relevant and so should be pressed into behaviour-guiding service. Unfortunately this strategy has to square up to a serious difficulty, since the putatively context-specifying second-order representations to which it appeals will need to have their own contextual relevance specified by other, third-order representations; these third-order representations will need to have their contextual relevance specified by fourth-order representations; and so on. The result is a computationally debilitating infinite regress (see e.g. Dreyfus 1992, pp.288-9).

What the foregoing considerations indicate is that cognitive science faces a *problem* of relevance, a problem of how to explain human sensitivity to relevance scientifically. So far, so good (or, perhaps, bad). What really concerns us here, however, is the way in which Heideggerian philosophy might be pressed into service to provide a diagnosis of, and a response to, the root cause of this difficulty.

Dreyfus's seminal, Heidegger-influenced analysis of why cognitive science has historically foundered in the face of the problem of relevance is well-known, and I do not intend to repeat all its details here. (Dreyfus gives a lively and compelling summary in his chapter for this

volume; for further details and discussion, see e.g. Dreyfus, 1990, 1992; Wheeler, 2005, 2008, 2010b; Cappuccio and Wheeler, 2010; Rietveld, this volume.) In any case, in the dispute with which we are concerned, it is agreed ground between Dreyfus, Rietveld and me (more accurately, Rietveld and I agree with Dreyfus) that the orthodox cognitive-scientific approach to the problem of relevance won't work and that an alternative approach is needed that is shaped by Heideggerian philosophy, or at least by the phenomenological tradition that Heidegger helped launch. That said, however, a key point in Dreyfus's critique of orthodox cognitive science (a point to be identified soon) will continue to be important in what follows, because Dreyfus's critical response to my version of Heideggerian cognitive science turns largely on the thought that I have failed to take on board what really matters in Heidegger's thinking, and that, because of this failure, the positive account of relevance-sensitivity that I end up giving does not put sufficient conceptual distance between itself and the account on offer from orthodox cognitive science, with predictably disastrous results. In short, as Dreyfus sees things, my lack of success in being sufficiently Heideggerian undermines not only my Heideggerian credentials but also my ability to offer a convincing response to the problem of relevance in any of its manifestations. That explains why Dreyfus describes my approach as advocating a "pseudo-Heideggerian AI" (Dreyfus, this volume, p.??) and concludes that I ultimately provide "no solution or dissolution of the [frame] problem" (Dreyfus, this volume, pp.??). If Dreyfus is right, I really have messed up.

To make a start on bringing the issues into proper view, I need to recap certain features of my version of Heideggerian cognitive science (Wheeler, 2005, 2008, 2010b). In Division 1 of *Being and Time*, Heidegger (1927) famously identifies three different modes of encountering entities: readiness-to-hand, un-readiness-to-hand, and presence-at-hand. Some aspects of this three-way distinction were described earlier, but now we need to add in more detail. (In truth, some of this added detail goes a little beyond what Heidegger himself actually says, but my intention at such moments is to mine the implications of Heidegger's own presentation, not to graft on alien ideas.) The mode of *readiness-to-hand* is instantiated in the online, hitch-free, skilled manipulation of equipment in certain sorts of context-specific practical tasks. In the domain of readiness-to-hand, the subject-object distinction, and thus representational consciousness, plays no experiential part. For example, while engaged in trouble-free text editing, the proficient typist will have no conscious recognition of the keyboard or the monitor, in the way that one would if one simply stood back and thought about them. Nor indeed will he have any experience of himself as a subject over and against his ongoing activity. Dreyfus (this volume) calls this kind of activity *absorbed coping*, and notes that it is regulated by (i) the human expert's capacity to sense deviations from a contextually determined optimal balance with her environment, and (ii) her ability to smoothly adapt her behaviour to improve her performance and thus reduce her sense of being out of balance.

When the kind of smooth adaptive accommodation that characterizes absorbed coping is disturbed by broken or malfunctioning equipment, discovered-to-be-missing equipment, or in-the-way equipment, our encounters with entities enter the new domain of *un-readiness-to-hand*, a domain in which entities are revealed phenomenologically as presenting context-specific problems to be solved. Skilled activity is consequently reconfigured as context-specific practical problem solving. When revealed as un-ready-to-hand, an entity will be experienced in terms of properties that are, for example, action-specific, egocentric and dependent on a particular context of activity. So, for example, the online, task-engaged skilled navigator, confronted by obstacles to be avoided, may represent the external environment as an egocentrically defined space in which those obstacles appear only as regions to be avoided, positioned in terms of roughly specified bearings relative to her own

body. In this way, un-readiness-to-hand involves the emergence of a proto-subject-object distinction that becomes more pronounced with increasing levels of disturbance, until eventually the entities under study are phenomenologically removed from the action-oriented settings of everyday equipmental practice and are thereby revealed as fully-fledged *present-at-hand* objects, that is, as the bearers of certain context-independent determinate or measurable properties (size in metres, weight in kilos etc.). The disclosure of entities as present-at-hand, which, as we have seen, is the characteristic mode of revealing of science, is accompanied by the full emergence of a subject over and against an objective universe, and thus of full-blown representational consciousness. So, for example, the disengaged reasoner, reflecting offline on a navigation problem, may represent the external environment by way of something like a Cartesian co-ordinate system in which the objective shape, orientation and positions of detected obstacles are plotted.

What kinds of causal mechanisms might cognitive science identify as enabling the modes of encounter just identified? Since our interest here is in the human capacity for fluid and flexible context-dependent sensitivity to relevance, and since the domain of presence-at-hand is a field of disengaged reflection on context-*independent* objects whose contextual relevance always remains to be determined (more on this later), I shall concentrate on the kind of contextually relevant skilled practical activity that is manifested in ready-to-hand and un-ready-to-hand scenarios, scenarios in which, as Heidegger (1927) illuminates things, human experience is of being *thrown* in a world that already matters to us, without there being any need for us to locate and specify that relevance. At this point we can use as our guide the (as I have argued) Heideggerian notion that there is an intelligible interplay between constitutive understanding and enabling understanding, and proceed to seek out candidate mechanisms that, in some interesting way, are structurally isomorphic counterparts to the phenomenologically identified structures described above (for a more careful justification of this appeal to structural isomorphisms, see Wheeler, 2005). The idea here is that we can neutralize the problem of relevance via a combination of Heideggerian phenomenology, through which we gain insight into the conditions of possibility of relevance-sensitive activity, and an empirically well-supported cognitive science which makes that phenomenological account intelligible from a mechanistic, causal-explanatory perspective.

It is plausible (according to me anyway) that both the ready-to-hand phenomenon of nonrepresentational sensitivity to contextually specific relevant factors and the un-ready-to-hand phenomenon of representationally mediated, context-specific problem solving may be enabled by mechanisms that I have described previously as *situated special purpose adaptive couplings* or, more catchily, as *thrown machines* (Wheeler, 2008). Thrown machines are closely coupled brain-body-environment mechanisms that, in effect, avoid the problem of having to determine what is relevant in a situation – and thus arguably solve the problem of relevance – because they are activated correctly only in the presence of the right, contextually relevant input. As I understand such thrown machines, they exhibit a kind of enabling-level *intrinsic context-sensitivity*. To illustrate this idea, here is an old but compelling example which I have used before (see e.g. Wheeler, 2004).

The evolutionary roboticists Harvey et al. (1994) set artificial evolution the task of designing a robot control system, comprising an artificial neural network and some rather basic visual receptors, that would enable a mobile agent to approach a triangle but not a rectangle in wildly varying lighting conditions. The evolutionary outcome was striking. Two visual receptors were positioned geometrically such that visual fixation on the oblique edge of the triangle would typically result in a pair of visual signals (i.e., receptor 1 = low, receptor 2 =

high) which was different from such pairs produced (almost) anywhere else in the arena. The robot would move in a straight line if the pair of visual signals was appropriate for fixation on the triangle, and in a rotational movement otherwise. Thus if the robot was fixated on the triangle, it would tend to move in a straight line towards it. Otherwise it would simply rotate until it did locate the triangle. Occasionally (and this explains the parenthetical 'almost' a few sentences ago) the robot would fixate erroneously on one edge of the rectangle, because, from certain angles, that edge would result in a qualitatively similar pair of visual signals being generated as would have been generated by the sloping edge of the triangle. Perturbed into straight line movement, the robot would begin to approach the rectangle. However, the looming rectangle would, unlike a looming triangle, produce a change in the relative values of the visual inputs (receptor 1 would be forced into a high state of activation), and the robot would be perturbed into a rotational movement. During this rotation, the robot would almost invariably refixate on the correct target, the triangle.

Here is a way of interpreting the activity of this thrown machine that, if we allow ourselves the obviously hazardous but potentially enlightening generalization to human activity, underlines its power as a model for an enabling account of absorbed coping. A mechanistic sensitivity to deviations from a contextually determined optimal balance with the environment (deviations from successful triangle tracking) elicit smooth adjustments to behaviour without the need for mediating representational control. Because of the specific nature of the adaptive dynamical coupling instantiated, there is never a stage of processing in which the mechanism 'pauses' to construct a map of its environment or produce a represented plan of action, no matter how partial, perspectival, rudimentary or short-lived one thinks the most minimal examples of such representational structures could be.

Although the thrown machine just described is nonrepresentational in character, the special-purpose and closely-coupled aspects of the mechanisms in question do not, in and of themselves, prevent the deployment of certain kinds of representational control strategies. So the presence of enabling-level representations is not necessarily in tension with intrinsic context-dependence. Put more positively, the point is that representational elements may partially underlie the phenomenon of context-dependent skilled practical problem solving in the domain of un-readiness-to-hand. (Supporting analyses for the claims made here may be found in Wheeler, 2005, 2008, 2010b; for an innovative empirical study that seeks to isolate the psychological signature of the *transition between* readiness-to-hand and unreadiness-to-hand, although admittedly without thereby endorsing the specific, representation-introducing aspect of my account, see Dotov et al., 2010.) Robots whose behaviour-generating inner elements arguably realize an analogue of the representational profile that, as Heidegger's analysis suggests, is phenomenologically distinctive of un-readiness-to-hand – in that they encode properties that are action-specific, egocentric and dependent on a particular context of activity – have been developed within (what I take to be) a 'thrown-machine' approach pursued by, among many others, Mataric (1991) and Franceschini et al. (1992). Indeed, the Franceschini et al. robot just referenced represents obstacles precisely as (to adapt my previous description) regions to be avoided, positioned in terms of roughly specified bearings relative to the robot's own body. Enabling-level representations with the kind of profile just sketched have been dubbed *action-oriented* representations (see e.g. Wheeler, 2005).

So much for the problem of relevance, one might think. At this juncture, however, there is a serious complication to be acknowledged, namely that the problem of relevance, which I have been treating as a unitary beast, actually has a *two-dimensional* structure (Wheeler, 2008, 2010b). Its *intra*-context dimension challenges us to say how a naturalistically

discharged system is able to achieve appropriate, flexible and fluid action *within* a context. It is this intra-context dimension of the problem for which thrown machines, as I have characterized them, provide a credible enabling understanding. But secondly, there is an *inter*-context dimension to the problem. This challenges us to say how a naturalistically discharged system is able to flexibly and fluidly switch between an open-ended sequence of contexts in a relevance-sensitive manner. For reasons that will become clear, I remain less confident regarding the enabling understanding of such context-switching, although, as I have argued previously (Wheeler, 2005, 2008, 2010a, b), part of the solution may well lie with mechanisms which realize a form of causation that Andy Clark once dubbed *continuous reciprocal causation* (Clark, 1997). Continuous reciprocal causation is causation that involves multiple simultaneous interactions and complex dynamic feedback loops, such that (a) the causal contribution of each systemic component partially determines, and is partially determined by, the causal contributions of large numbers of other systemic components, and (b) those contributions may change radically over time. This species of causation plausibly bestows on a mechanistic system a certain kind of large-scale holistic flexibility, a flexibility that seems to be ripe to account, in part, for the fluid context-switching highlighted by the inter-context frame problem. More on this, with an example, in a moment.

That completes my brief recap of the salient aspects of my version of Heideggerian cognitive science. So now, what precisely might be wrong with it?

5. Representations and Background Coping

At the heart of the disagreement between Dreyfus and me are two intertwined issues, namely (i) the nature, status and consequences of *representational explanation* in the cognitive science of relevance-sensitivity, and (ii) the existence, character and contribution to human relevance-sensitivity of a phenomenon that Dreyfus calls *background coping*.

Dreyfus (this volume, pp.??) deplores what he calls my "cognitivist misreading" of Heidegger for its appeal to representations and (as Dreyfus sees it) its associated assumption that skilled human activity can be treated exhaustively as problem-solving. Although I might take issue with the latter claim – since I don't treat absorbed coping as problem-solving (see above) – that is a minor interpretative quibble. The real action here concerns my (as I see it) Heideggerian thought that action-oriented representations figure in the kind of relevance-sensitivity that characterizes practical problem-solving in the domain of the un-ready-to-hand. It is important to note that this is not (or not only) an argument over who gets to wear the Heidegger t-shirt. The full scope of Dreyfus's objection is that there are good Heideggerian reasons for thinking that the problem of relevance is, at least in part, *an artefact of representationalism*. As he puts it, "for Heidegger, all representational accounts are part of the problem" (Dreyfus, this volume, p.??). Perhaps the central consideration in support of this claim is Heidegger's somewhat sketchy treatment of what he calls *value-predicates* (Heidegger, 1927, p.132; for discussion, see Dreyfus, 1990; Wheeler, 2005). In effect, value-predicates are representations of context-dependent features. And Heidegger argues that adding such representations to context-independent, present-at-hand primitives (e.g. to raw sense data or, to give the argument a more contemporary tone, to light-intensity gradients at the retina) can never be the ultimate source of relevance, since each such representation is itself a present-at-hand structure that needs to have *its* contextual relevance determined. From a Heideggerian perspective, this is the fundamental root of the cognitively incapacitating infinite regress that, as we saw earlier, arguably thwarts the orthodox cognitive-scientific approach to the problem of relevance.

Given all this, Dreyfus's criticism might seem to be that because I fail to follow Heidegger's recommendation, which (the treatment of value-predicates might suggest) is to do away with representations altogether, my so-called Heideggerian cognitive science will inherit something akin to the difficulties that have plagued orthodox cognitive science. There is certainly textual evidence for this interpretation of Dreyfus. For example, he writes that "Wheeler's approach… by introducing flexible action-oriented representations, like any representational approach has to face the frame problem head on" (Dreyfus, this volume, pp.??). If this were Dreyfus's settled position, however, it would seemingly be refuted by the fact that, as we have seen, there are examples of representation-involving thrown machines that, because they operate on the basis of close coupling and special-purpose activation, don't succumb to the problem of relevance. Crucially, the action-oriented representations that figure in such machines are not designed to perform the function of specifying context (context is *intrinsic* to the operative mechanisms; see above). This suggests that the structure of the orthodox problem of relevance is not an artefact of representationalism *per se*, but rather an artefact of the idea that context can be representationally specified. So Dreyfus's criticism, as I am currently understanding it, looks to be unfounded. On reflection, however, things can't be quite that simple. That is, and despite the occasional rhetorical flourish, Dreyfus cannot hold (and cannot hold that Heidegger would hold) that the mere introduction of representations anywhere into an account of what constitutes and enables relevance-sensitivity results in that account becoming impaled on the problem of relevance. And that's because Dreyfus himself allows that *un-readiness to hand and the richly relevance-sensitive skilled activity that accompanies it have a representational signature* (see e.g. Dreyfus, this volume, p.??). Dreyfus's critical point, then, is more subtle. It is that the representational signature in question plays no interesting role in accounting for relevance-sensitivity. This deflationary judgment turns on a Heideggerian claim to which, Dreyfus suggests, I am all but blind. That claim is that "all coping, including unready-to-hand coping, takes place on the background of [a] basic nonrepresentational, holistic, absorbed, kind of intentionality, which Heidegger calls being-in-the-world" (Dreyfus, this volume, pp.??).

What, then, is this all-important phenomenon of the background? Here is a brief sketch (for a more detailed treatment, see Cappuccio and Wheeler, forthcoming). In the Heideggerian framework, the significance of the world into which human activity is thrown is neither fully represented within the practical knowledge that enables Dasein to negotiate situations, nor is it fully representable in some ideal theoretical register. This is because that structure encompasses a vast and indeterminate web of implicit preconditions for sense making, "an unexplicated horizon" or *background*, providing "the vantage point from out of which" every experience matters to one in certain way (Taylor, 1993, p.325). Our epistemic engagement with the background thus amounts to a knowing how to navigate one's way around those psychological, social and cultural structures of one's world that are implicitly presupposed by, rather than on open display in, concrete examples of skilled activity. Dreyfus (this volume) calls our familiarity with, and our smooth navigation of, these holistic and unrepresented patterns of significance *background coping*. Background coping may thus be closely allied with the Heideggerian sense of belonging to the world identified by Ratcliffe (see above).

So, phenomenological analysis reveals background coping to be a nonrepresentational, transcendental feature of human existence that, as a matter of ontology, underlies both nonrepresentational absorbed coping and representational practical problem solving. From the vantage point of a Heideggerian cognitive science, what remains to be given is an

enabling explanation of this fundamental capacity. When Dreyfus (this volume) turns his attention to this outstanding issue, he cites with approval a scientific research programme that (as I see it) emphasizes the very phenomenon of continuous reciprocal causation whose potential for underpinning large-scale holistic flexibility I highlighted earlier. That programme is the neurodynamical framework developed by Freeman (2000). (For remarks by Freeman which all but prefigure the language of continuous reciprocal causation, see Freeman, 2000, p.22; quoted by Dreyfus, this volume, pp.??). In Freeman's research (more on which in a moment), the brain is conceived as, broadly speaking, a nonrepresentational dynamical system primed by past experience to actively pick up and enrich significance, a system whose constantly shifting attractor landscape causally explains how newly encountered significances may interact with existing patterns of inner organization to create new global structures for interpreting and responding to stimuli. With this candidate for an enabling understanding of background coping in place, the Dreyfusian version of Heideggerian cognitive science comes fully into view.

> Showing in detail how the representational un-ready-to-hand in all its forms depends upon a background of holistic, nonrepresentational coping is exactly the Heideggerian project and would, indeed, be the most important contribution that Heideggerian AI could make to Cognitive Science. Indeed, a Heideggerian Cognitive Science would require working out an ontology, phenomenology, and brain model, that denies a basic role to any sort of representation – even action oriented ones – and defends a dynamical model… that gives a primordial place to equilibrium and in general to rich coupling. (Dreyfus, this volume, pp.??)

One way for me to respond to Dreyfus's critique would be for me to argue directly that there is a representational dimension to the way in which we navigate the background. I think that such an argument can be made, based on the way in which certain embodied routines, such as the preparatory exercises sometimes executed by sporting experts, reconfigure the background to promote future behavioural success while being interpretable as a species of action-oriented representation (Cappuccio and Wheeler, forthcoming). My strategy here, however, will be to argue that, when all the cards are on the table, neither Dreyfus nor I is currently in a position to offer a fully compelling Heideggerian cognitive science of relevance-sensitivity, but that my account fares better than Dreyfus's and thus offers a reasonable platform for future thought. As we shall see, a crucial step in the reasoning by which I reach this conclusion will involve an examination of Rietveld's objections to my views.

Recall the inter-context problem of relevance, the problem of specifying the causal mechanisms that underlie flexible and fluid context-switching. According to Dreyfus's analysis, such context-switching is encompassed by our capacity for background coping, a capacity distinguished phenomenologically by a broad implicit sensitivity to holistically structured global relevance and explained causally by something like Freeman-style dynamics. This explanatory package is on display in the following passage, in which Dreyfus draws on Freeman's work to address the issue of context-switching. Dreyfus writes:

> If Freeman is right… our sense of other potentially relevant familiar situations on the horizon of the current situation [i.e., our readiness and capacity for context-switching], might well be correlated with the fact that brain activity is not simply in one attractor basin at a time but is influenced by other attractor basins in the same landscape, as well as by other attractor landscapes which under what have

previously been experienced as relevant conditions are ready to draw current brain activity into themselves. According to Freeman, what makes us open to the horizonal influence of other attractors is that the whole system of attractor landscapes collapses and is rebuilt with each new rabbit sniff [Freeman has worked extensively on rabbit olfaction], or in our case, presumably with each shift in our attention. And after each collapse, a new landscape may be formed on the basis of new significant stimuli – a landscape in which, thanks to past experiences, a different attractor is active. (Dreyfus, this volume, pp.??)

There are reasons to think that Dreyfus's appeal to shifting attractor landscapes falls short of what is required (Wheeler, 2010a, b). The first thing to notice is that it remains unclear from the text in question whether the crucial reconfiguration of the neural attractor landscape is supposed to be (i) caused by the attentional shift (as might be suggested by the parallel with the rabbit sniff and the talk of a new landscape being formed "on the basis of new significant stimuli") or (ii) the causal basis of the attentional shift (as might be suggested by the thought that the attractors in the landscape determine what we attend to). Either way, there is a worry. If (i) is the correct interpretation, then the shift in attention itself remains unexplained. But at least sometimes that shift in attention is presumably governed by, and thus presupposes, a grip on the way in which contexts of activity are changing. To that extent, then, Dreyfus' suggestion begs the question. On the other hand, if (ii) is the correct interpretation, then it seems that we are still owed an explanation of how it is that, out of all the attractors in the pre-transition landscape that have been significant in the past, and that might have become active, it is the relevant one that is ultimately selected. Once again, it seems, the key question is being begged.

If the foregoing worries are justified, Dreyfus's favoured candidate for an enabling understanding of fluid and flexible context-switching is, at best, incomplete. Of course, the wily reader will be champing at the bit right now to point out that if Freeman-style neurodynamics are (as I have suggested) a species of continuous reciprocal causation, and if (as I have indicated) my own preferred candidate for a mechanism underlying context-switching is one that instantiates continuous reciprocal causation, then Dreyfus and I look to be in the same (leaky) boat. There is undoubtedly something right about this. The fact that a machine may flexibly and holistically reconfigure itself on the basis of continuous reciprocal causation among its elements does not guarantee that the behaviours generated by that machine will remain contextually relevant. All that is assured is that the machine supports the kind of flexibility that, when harnessed appropriately (i.e. in relevance-sensitive ways), may help to generate fluid context-switching. In other words, although Dreyfus and I may be in possession of part of the story about how fluid and flexible context-switching is causally enabled, that story remains radically unfinished (Wheeler, 2010a).

Despite this common shortfall, however, the situation is arguably worse for Dreyfus than it is for me. For although, in the quoted passage, Dreyfus is focussed on the way in which background coping grounds context-switching, his position, as I understand it, is that it is at root the same set of complex, dynamical, nonrepresentational structures and processes that materially enables both intra-context sensitivity to relevance (whether absorbed or disturbed) and context-switching. This is seemingly a straightforward consequence of Dreyfus's claim (see above) that all coping, whether ready-to-hand or un-ready-to-hand, depends on background coping, plus his favoured account of the causal mechanisms underlying the latter. As we have just seen, though, there is an acute worry hereabouts, namely that any appeal to the kinds of structures and processes highlighted by Dreyfus presupposes a further account of

mechanistic relevance-sensitivity to explain precisely how those elements are harnessed in appropriate ways. As far as I can tell, this worry concerning the explanatory power of the dynamical systems in question applies just as much to the intra-context case as it does to the inter-context case. And it is in the light of this observation that my view emerges with the advantage. For although I share the highlighted problem in the inter-context case, I do not share it in the intra-context case. My claim, pace Dreyfus, is that intra-context sensitivity to relevance is materially enabled by (what I called earlier) thrown machines, special-purpose and closely coupled brain-body-environment mechanisms. This claim is not undermined by the considerations on the table.

We now have a new dispute in view, one that concerns the relative scope and power of continuous reciprocal dynamics and thrown machines within a Heideggerian cognitive science. Can we resolve this new dispute? Rietveld thinks we can. Unfortunately for me, however, he thinks that the resolution comes out in Dreyfus's favour. To bring my self-interested defence to a close, then, I shall attempt to resist Rietveld's arguments.

6. The Dimensionality of the Problem of Relevance

Rietveld (this volume) argues that my distinction between intra-context sensitivity to relevance and inter-context-sensitivity to relevance is ultimately unsustainable, on the grounds that one can speak of genuine sensitivity to relevance only where there exists the intrinsic possibility of context-switching, that is, where human activity remains essentially sensitive to the holistic structure of the background. In Rietveld's own words, the "distinction does not make much sense for understanding normal behavior, because in cases of human (and animal) online intelligence, sensitivity to real relevance is not a feature of responsiveness to some limited part of the context (as "intra-contextual" sensitivity to relevance would imply) but of responsiveness to the individual's full particular and complex situation" (Rietveld, this volume, p.??). If Rietveld's reasoning is sound, it would provide support for a Dreyfus-style account of the relationship between background coping and absorbed/disturbed coping, since there would be no precise boundary separating the context-specific forms of coping from holistic background coping, but only local ways for the latter to be revealed within the former through the ongoing evolution of the shifting attractor landscapes that characterize the underlying neurodynamics. Correlatively, there would be something seriously misguided about my attempt to carve out a causal domain in which thrown machines operate in such a way that they are somehow encapsulated from holistic significance. Thus, Rietveld concludes, "[s]ince Wheeler has already agreed with Dreyfus that continuous reciprocal causation… underlies contextual sensitivity ("inter-context" switching) on broader scales, [Wheeler] should accept that in the end [such causation] turns out to be the mechanical basis of *all* online sensitivity to real context-dependent relevance, not just of inter-context sensitivity" (Rietveld, this volume, p.??).

Rietveld's case against the contested distinction involves two converging lines of argument. According to the first, if we use neuropsychology to determine the location of the most plausible candidates for thrown machines in humans (should any such mechanisms exist), it turns out that they will be found in the lateral premotor cortex. However, it also transpires that the successful functioning of those machines depends on prior (i.e., pre-sensing) activity in the medial premotor system. The experimental data suggests that this prior activity tunes the candidate mechanisms in the lateral premotor cortex to what is currently relevant to the individual concerned, by orienting those mechanisms' sensory expectations (see Rietveld, this volume, for supporting data and references). The key idea, I think, is this: since

background-sensitive global tuning determines the local sensory-sensitivities of the candidate special-purpose thrown machines, one cannot hold onto the view that those machines are functionally isolated from the background, in the way that the distinction between inter-context and intra-context sensitivity to relevance would suggest.

In considering this argument, let me begin by recording a concern about the way in which Rietveld sets up his example. Bemoaning the fact that I am "not explicit about the specific neural mechanisms that could underlie human "intra-context" sensitivity to relevance" (Rietveld, this volume, p.??), he locks on to the lateral premotor cortex as a likely site for thrown machines in human beings, should any such machines exist. However, I am not inclined to think of the paradigm cases of thrown machines as purely neural routines, but rather as distributed partnerships that incorporate neural routines, non-neural bodily elements, and environmental factors. Consider, for example, the evolved robotic mechanism for triangle-square discrimination discussed earlier, in which significant contributions to adaptive success are made not only by the activation profile of the inner neural network, but also by the geometric arrangement of the robot's peripheral sensors and by the specific structure of the environment (the mechanism would presumably fail to track only triangles in environments populated by sloping edges that weren't parts of triangles). This makes me suspicious that Rietveld's purely neural routines couldn't really count as complete thrown machines at all. Waiving this concern, however, the mere fact that a candidate thrown machine depends on prior processing elsewhere in the system to orient its sensitivities before it becomes active does not, as far as I can tell, undermine the encapsulated status of that mechanism *during its processing cycle*. This way of thinking about thrown machines becomes clearer, I think, once it is noted that such machines may emerge as transient phases of modularity in a system that, at other times, exhibits continuous reciprocal causation (for an example, see Wheeler, 2005, p.265). In such dynamically shifting mechanisms, the pre-modular phases of holistic reorganization presumably play a role in establishing the conditions for the subsequent online operation of the emergent thrown machines. But if this picture is right, then the existence of the kind of prior processing on which Rietveld hangs his argument does not destabilize the distinction between inter-context and intra-context relevance-sensitivity. Rather, it usefully highlights an aspect of the relationship *between* the two phases of operation.

Rietveld's second line of argument depends on the thought that a certain psychopathology, so-called *utilization behaviour* (UB; see, e.g., Lhermitte, 1983), tells us what human activity would be like, if it were in fact underpinned by thrown machines. UB sufferers are sensitive to the local affordances provided by the immediate environment in a way that is insensitive to the distinction between socially appropriate and socially inappropriate actions. As a result, although the sufferer's behavioural responses remain "technically adequate" (Rietveld, this volume, p.??), contextually inappropriate actions are not inhibited, even though broader background knowledge should discourage them. For example, a UB sufferer may enthusiastically launch into the act of making a bed, even if this bed is in someone else's house. Rietveld's diagnosis is that UB results from a disruption of the normal tendency of human agents to situate the local context of action within its holistic background. As he puts it, "immediate responsiveness to affordances has lost its sensitivity to the full particular situation" (Rietveld, this volume, p.??). And this, he argues, is how all human activity would be if it were underpinned by thrown machines. "Patients with UB manifest what Wheeler would call "intra-context" sensitivity to relevance (they act in technically adequate ways). Yet, given their full particular situation (or context) their actions are inappropriate and we can

characterize their behavior as form of responsiveness to irrelevant affordances" (Rietveld, this volume, p.??).

What strikes me as problematic about this argument is Rietveld's claim that UB patients realize a properly working thrown machine. Indeed, as Rietveld rightly notes, technical adequacy is not sufficient for relevance-sensitivity. To explain the missing relevance-sensitivity of such merely technically adequate behaviour, I am inclined to conclude not, with Rietveld, that thrown machines are incapable of genuine relevance-sensitivity, but that UB patients either have a faulty thrown machine (in effect, one that isn't thrown) or have a breakdown at the interface between their thrown machines and whatever prior processing might be in place to orient the sensitivities of those machines (see above). Ironically, given Rietveld's intentions, one might develop the latter possibility via the suggestion that the behaviour of the UB patient is explained by the fact that the separation between local context and background has collapsed in such a way that the background is now entirely explicated by the actual focus of interest, such that no other contexts are possible (Cappuccio and Wheeler, forthcoming). It is as if the sensitivities of his thrown machine have been set up in such a way that the bed is the totality of the world with which he must deal.

To this sort of response, Rietveld will object that to make good on the idea that thrown machines meet the adequacy condition on relevance-sensitivity (which, roughly, is that they must exhibit the kinds of sensitivities that UB patients lack), I would need to expand what counts as intra-contextually relevant. For example, as he notes, I might claim that the bed-making context should include not only the made or unmade state of the bed, but also whether or not the bed is in someone else's house. But this strategy, he argues, will end up collapsing the very distinction that I am trying to shore up, since once one starts to enlarge the set of context-specifying factors, it seems that there is no properly motivated principle that could lead us to stop short of fading into the background itself. Thus he writes: "[p]ulling this string of broadening the "intra-context" does not help Wheeler, because it amounts to requiring the kind of global holism that "intra-context sensitivity to relevance" by definition lacks", such that "the individual ends up coping with exactly the same full and particular (or "inter") context that Wheeler had wanted to reserve for context switching" (Rietveld, this volume, p.??).

This is, I admit, a genuine risk, one that casts a spotlight on the Herculean theoretical challenge of giving a theory of how our cognitive systems determine where the boundaries of contexts lie. It seems to me, however, that at this juncture we confront a choice. Either we remain committed to the distinction between intra-context and inter-context sensitivity to relevance, and so humbly accept the challenge to give a theory of context determination, or we follow Dreyfus and Rietveld in rejecting that distinction, and so confront the fact that, as I argued earlier, our Heideggerian cognitive science of relevance-sensitivity in general, and not just of context-switching, has barely begun. At present I am backing the first of these horses, but that is a bet whose fate will be decided by the ongoing intelligible interplay between phenomenological philosophy and empirical cognitive science that characterizes the still-unfolding and intellectually exciting story that is Heideggerian cognitive science.

References

Cappuccio, M. and Wheeler, M. (2010), 'When the Twain Meet: Could the Study of Mind be a Meeting of Minds?', in J. Reynolds, E. Mares, J. Williams and J. Chase (eds.), *Postanalytic and Metacontinental: Crossing Philosophical Divides*. London: Continuum.

Cappuccio, M. and Wheeler, M. (forthcoming), 'Ground-Level Intelligence: Action-Oriented Representation and the Dynamics of the Background', in Z. Radman (ed.), *Knowing without Thinking: Mind, Action, Cognition and the Phenomenon of the Background*.

Clark, A. (1997), *Being There: Putting Brain, Body, And World Together Again*, MIT Press, Cambridge, Mass..

De Preester, H. (2008), 'From Ego to Alter Ego: Husserl, Merleau-Ponty and a Layered Approach to Intersubjectivity'. *Phenomenology and the Cognitive Sciences*, 7 (1): 133–142.

Dotov, D.G., Nie, L., Chemero, A. (2010), 'A Demonstration of the Transition from Ready-to-Hand to Unready-to-Hand'. *PLoS ONE* 5(3): e9433. doi:10.1371/journal.pone.0009433

Dreyfus, H. L. (1990), *Being-in-the-World: A Commentary on Heidegger's Being and Time*, Division I, Cambridge, Mass.: MIT Press.

Dreyfus, H. L., (1992), *What Computers Still Can't Do: A Critique of Artificial Reason*, Cambridge, Mass.: MIT Press.

Dreyfus, H. L. (This volume), 'Why Heideggerian AI Failed and how Fixing it would Require Making it more Heideggerian'. Published previously in P. Husbands, O. Holland and M. Wheeler (eds.) (2008), *The Mechanical Mind in History*. Cambridge, Mass.: MIT Press, pp. 331–71. A shortened version appears under the same title in *Philosophical Psychology*, 20 (2) (2007), 247–68. A further version appears under the same title in *Artificial Intelligence*, 171 (2007), 1137–60.

Elton, M. (2003), *Daniel Dennett: Reconciling Science and our Self-Conception*. Cambridge: Polity Press.

Fodor, J. A. (1983), *The Modularity of Mind*, Cambridge, Mass.: MIT Press.

Franceschini, N., Pichon, J. M., and Blanes, C. (1992), 'From Insect Vision to Robot Vision, *Philosophical Transactions of the Royal Society, series B*, 337, 283-94.

Gallagher, S. (2003), 'Phenomenology and Experimental Design', *Journal of Consciousness Studies*, 10 (9-10), 85-99.

Gallagher, S. and Zahavi, D. (2008), *The Phenomenological Mind: an introduction to Philosophy of Mind and Cognitive Science*. London and New York: Routledge.

Haddock, A. (2009), 'McDowell, Transcendental Philosophy, and Naturalism', *Philosophical Topics*, 37 (1), 63-76.

Harvey, I., Husbands, P., and Cliff, D. (1994), 'Seeing the Light: Artificial Evolution, Real Vision', in *From Animals To Animats 3: Proceedings Of The Third International Conference On Simulation Of Adaptive Behavior*, D. Cliff, P. Husbands, J.-A., Meyer and S.W. Wilson (eds.). Cambridge, MA.: MIT Press, pp.392-401.

Heidegger, M. (1927), *Being and Time*, trans. J. Macquarrie and E. Robinson. Oxford: Basil Blackwell, 1962.

Heidegger, M. (1954), 'The Question Concerning Technology', translated by W. Lovitt with revisions by D. F. Krell, in D. F. Krell (ed.) *Martin Heidegger: Basic Writings*, revised and expanded edition. London: Routledge, 1993, pp.311-41.

Kelly, S. (2000), 'Grasping at Straws: Motor Intentionality and the Cognitive Science of Skilled Behaviour', in M. Wrathall and J. Malpas (eds.) (2000), *Heidegger, Coping and Cognitive Science: Essays in Honor of Hubert L. Dreyfus*, vol. 2. Cambridge, Mass.: MIT Press.

Kuhn, T. (1962), *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Lakatos, I. (1977), *The Methodology of Scientific Research Programmes*. Cambridge: Cambridge University Press.

Lhermitte, F. (1983) '"Utilization Behaviour" and its Relation to Lesions of the Frontal Lobes', *Brain* 106, 237-255.

Mataric, M. (1991), 'Navigating with a Rat Brain: a Neurobiologically Inspired Model for Robot Spatial Representation', in J.-A. Meyer and S. Wilson (eds.) *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, Cambridge, Mass.: MIT Press, pp.169-75.

McDowell, J. (1994), 'The Content of Perceptual Experience'. *The Philosophical Quarterly* 44 (175), 190-205.

McDowell, J. (1996), *Mind and World*, 2nd edition. Cambridge, Mass.: Harvard University Press. 1st edition published 1994.

McDowell, J. (1999), 'Sellars on Perceptual Experience', in his *Having the World in View: Essays on Kant, Hegel and Sellars*. Cambridge, Mass.: Harvard University Press.

Price, H. (2004), 'Naturalism without Representationalism', in D. Macarthur and M. de Caro (eds.), *Naturalism in Question*. Cambridge, Mass.: Harvard University Press.

Ratcliffe, M. (This volume), 'There can be no Cognitive Science of Dasein'.

Rehberg, A. (This volume), 'Heidegger and Cognitive Science – Aporetic Reflections'.

Rietveld, E. (This volume), 'Context-Switching and Responsiveness to Real Relevance'.

Rizzolatti, G. and Sinigaglia C. (2007), 'Mirror Neurons and Motor Intentionality'. *Functional Neurology*, 22(4), 205-210.

Sheehan, T. (2001), 'A Paradigm Shift in Heidegger Research', *Continental Philosophy Review*, 32(2), 1-20.

Taylor, C. (1993) 'Engaged Agency and Background in Heidegger' in C. B. Guignon (ed.) *The Cambridge Companion to Heidegger*. New York: Cambridge University Press, pp.317-36.

Wheeler, M. (2004), 'Is Language the Ultimate Artefact?', *Language Sciences*, 26(6), 693-715.

Wheeler, M. (2005), *Reconstructing The Cognitive World: The Next Step*. Cambridge, Mass.: MIT Press.

Wheeler, M. (2008), 'Cognition in Context: Phenomenology, Situated Robotics and the Frame Problem', *International Journal of Philosophical Studies*, 16(3), 323-49, special issue on Situated Cognition: Perspectives from Phenomenology and Science, M. Ratcliffe and S. Gallagher (eds.).

Wheeler, M. (2010a), 'Plastic Machines: Behavioural Diversity and the Turing Test', *Kybernetes*, 39 (3), 466-80.

Wheeler, M. (2010b), 'The Problem of Representation', in S. Gallagher and D. Schmicking (eds.), *Handbook of Phenomenology and Cognitive Science*. Dordrecht: Springer, pp.319-336.

Wheeler, M. (Forthcoming), 'Martin Heidegger', *Stanford Encyclopedia of Philosophy*, http://plato.stanford.edu/.

Wheeler, M. and di Paolo, E. (Forthcoming), 'Existentialism and Cognitive Science', in J. Reynolds, A. Woodward, and F. Joseph (eds.), *Continuum Companion to Existentialism*. Continuum, London.

Wolin, R. (ed.), *The Heidegger Controversy: a Critical Reader*, Cambridge, Mass.: MIT Press, 1993.