THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head

Citation for published version:

Clark, A 2009, 'Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head', *Mind*, vol. 118, no. 472, pp. 963-993. <https://doi.org/10.1093/mind/fzp110>

Digital Object Identifier (DOI):

[10.1093/mind/fzp110](https://doi.org/10.1093/mind/fzp110)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Mind

Publisher Rights Statement:

This is a pre-copy-editing, author-produced PDF of an article accepted for publication in *Mind* following peer review. The definitive publisher-authenticated version, Clark, A. 1 (2009) "Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head", *Mind* 118 (472) p. 963-993, is available online at: <http://mind.oxfordjournals.org/content/118/472/963>

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head

Andy Clark

Abstract

Is consciousness all in the head, or might the minimal physical substrate for some forms of conscious experience include goings on in the (rest of the) body and the world? Such a view might be dubbed (by analogy with Clark and Chalmers (1998) claims concerning ‘the extended mind’) ‘the extended conscious mind’. In this paper I review a variety of arguments for the extended conscious mind, and find them flawed. Arguments for extended cognition, I conclude, do not generalize to arguments for an extended conscious mind.

1. A Radical Response

Consciousness is puzzling. So puzzling, indeed, as to lead Jerry Fodor to assert that:

[We don’t know], even to a first glimmer, how a brain (or anything else that is physical) could manage to be a locus of conscious experience. This ... is, surely, among the ultimate metaphysical mysteries; don’t bet on anybody ever solving it. (Fodor 1998, p. 83)

Colin McGinn is no more optimistic, claiming that:

It is not that we know what would explain consciousness but are having trouble finding the evidence to select one explanation over the others; rather, we have no idea what an explanation of consciousness would even look like. (McGinn 1991, p. 61)

In recent years, however, some theorists (both within philosophy and cognitive science) have begun to suspect that our puzzlement may have a rather unexpected source. Perhaps, they suggest, we are simply *looking in the wrong place* for the mechanisms whose action explains the conscious mind. In particular, the suspicion is growing that the explanation of our qualitative mental life (the elusive 'what-it-is-likeness'ⁱ that seems to characterize a subject's experience of a certain kind of redness, of a certain voice, or of a pain in her stomach) has been hampered by a kind of blinkered vision that too firmly divides the neural from the extra-neural (gross bodily and environmental) realms. Such a view is prominent in, for example, recent collaborative work by the neuroscientist Diego Cosmelli and the philosopher Evan Thompson. Their radical but intriguing suggestion is that switching our explanatory focus from the brain to the embodied and situated brain (the-brain-plus-other-physical-stuff) will help turn the mystery into a (mere) puzzle, since 'the processes crucial for consciousness are not confined to the brain but include the body embedded in the environment' (Cosmelli and Thompson, forthcoming, ms p.1).

Nor are they a lone voice. Similar claims can be found in Varela, Thompson, and Rosch (1991), Van Gelder and Port (1995), Hurley (1998), Thompson and Varela (2001), Hurley and Noë (2003), Noë and Thompson (2004a, b), and Noë (2008). Driving this emerging view, it seems to me, are arguments and considerations drawn from three related, and increasingly popular, areas: the depiction of mind as essentially dynamic and 'process-like' in nature (Kelso, 1995, 2002), the development of the 'enactive' model of perception (O'Regan and Noë, 2001; Noë, 2004), and arguments for the 'extended mind' (Clark and Chalmers, 1998; Clark, 2007).

The attraction of a situated approach to consciousness are hard to underestimate. As one (skeptical) commentator recently put it:

Consciousness is trendy ... embodiment and situated cognition are also trendy ... both topics are exciting and being exciting is an additive property. An embodied/situated theory of consciousness is the philosophical equivalent of a blockbuster ...'. (Prinz 2009, p. 419)

In what follows, I assess the arguments (more cautiously, those arguments rooted in dynamic and enactive conceptions of mind) meant to favour an extended (more-than-neural) physical basis for conscious experience. I

ⁱ For a thorough examination of this important but elusive notion, see Chalmers (1995). For a critical view, see Dennett (1991).

conclude that most of these arguments visibly fail. One argument (concerning complex dynamic entanglement) does better, but seems threatened by some uncooperative empirical facts. I conclude that as things stand, there are no good reasons (of a dynamical, enactive, stripe) to endorse the vision of an extended conscious mind.

2. The Extended Mind

Clark and Chalmers (1998), offer an argument designed to show that cognition and mental states may, in human agents, sometimes depend on material mechanisms that extend beyond the boundaries of the brain and central nervous system. Some of these mechanisms may involve gross bodily structures (the cognitive role of gesture is a possible real-world case (Clark, 2007, 2008)), while others may involve extra-bodily resources such as computers and even good old-fashioned notebooks. Believers in the extended mindⁱⁱ (henceforth, EM) allow, importantly, for vast swathes of content-shaping causal contact between brain, (the rest of the) body, and world. Body, world, and action quite plainly shape and structure the contents of mind and of experience. But the EM hypothesis denies that this exhausts the potential role of well-matched non-neural resources in determining at least some of the mental and cognitive states and processes of an individual agent. Instead, in some cases, body and world are said to play what Hurley (1998), usefully describes as a non-instrumental role: they form part of the actual machinery of mind, acting as the so-called ‘physical vehicles’ (more on which below) of certain mental states and contents. Thus, the fact that moving the head brings something new into view and thus alters our mental states is not in itself surprising, and it should not incline us to endorse any claims about bodily extended vehicles for the new contents thus made available. For those new contents may still be encoded or carried by the altered patterns of purely neural activity caused by turning the head. Such merely causal dependence (of the mental upon the bodily or environmental) is what Hurley means by ‘instrumental dependence’. By contrast, the ensuing neural activity might be thought to be sufficient for the agent to come to believe (for example) that there is now a pink elephant in the room. This is the kind of non-instrumental involvement of material structure (sometimes called ‘constitutive involvement’ (Block, 2005)) that is at issue in debates concerning cognitive extension. EM claims that this kind of constitutive involvement, commonly (though not universally) granted to the neural activities underlying mental states and processes, can (under some

ⁱⁱ They include, though with differing perspectives and emphases, Haugeland (1998), Dennett (1996), Wilson (2004), Hutchins (1995, in press), Wheeler (2005, in press), Menary (2007), Tribble (2005), Sutton (2002), Noë (2004), Rowlands (2003, 2006), and Hurley (1998).

circumstances) also characterize the role of key *extra-neural* material structures. Examples might include the well-integrated use of a notebook as a non-neural data store (Clark and Chalmers, 1998), or those uses of hand and arm gestures that (it has more recently been suggested) are active elements in the construction of trains of thought and reason (Goldin-Meadow, 2003; McNeill, 2005; Clark, 2007).

It is important to noticeⁱⁱⁱ that all the claims at issue in this debate are claims about what are sometimes described as the ‘vehicles’ of mental states and processes. They are not claims about their contents or about how mental states present the world to the subject. Thus suppose you have a thought whose content is that the beach is distant. The thought is about extra-neural states of affairs. But the circuitry whose whirrings and grindings realize the thinking is, rather plausibly, entirely local. It consists, at least on most contemporary materialist models, in patterns of activity in neural populations. These neural populations, thus active, may be said to provide the local material vehicles of the mental content (for some discussions of the vehicle/content distinction, and its importance for theories such as EM (Hurley, 2003, forthcoming).

The EM thesis aims to put pressure on this received image of the ‘brain-bound’ (Clark 2008, p. xxv-xxix) nature of the material vehicles of mind and cognition. EM is thus the claim that the local material vehicles of some aspects of human cognition may, at times, be spread across brain, body, and world. This also (and we will pursue this strand in much more detail below) emerges as the claim that the local operations that realize some human cognizings include (possibly quite complex) tangles of feedback, feedforward and feedaround loops that promiscuously criss-cross the boundaries of brain, body and world.

The key case of human mental extension treated in Clark and Chalmers (1998), concerned the potential role of extra-neural structures (such as a fluently deployed, constantly available notebook) as apt for inclusion among the material vehicles of some of an agent’s non-occurrent (dispositional or ‘standing’) beliefs. Otherwise put, inscriptions in the notebook figure as part of the physical supervenience base for certain standing beliefs of the agent. Very roughly the argument was that for the normally ecologically situated brain it often does not matter whether information is stored in the head or left out in the world, just so long as the right information is retrieved or reconstructed at the right time, so as to govern actions in much the way we normally associate with antecedently holding the standing belief in question (Clark, 2008).

ⁱⁱⁱ Thanks to Keith Allen for reminding me of this.

The arguments in Clark and Chalmers (1998), and more recently in Clark (2008), thus aimed to show that whether a human individual counts as believing some fact is not simply a matter of whether it is stored in your biological memory (imagine if it was, but was totally inaccessible to recall), so much as a matter of whether the information thus encoded is appropriately poised for the guidance of behaviour. More generally, skin and skull do not, to use another phrase from Hurley (1998), form a ‘magical membrane’ within which (and only within which) real cognitive mechanisms (genuine vehicles of mental content) can be found. Clark and Chalmers thus defended a form of ‘active externalism’ in which (quite unlike standard philosophical externalisms about content) the relevant parts of the world are right there in the loop, active in the here-and-now. Such active externalism is easily distinguished from the more standard varieties of externalism^{iv} (such as those suggested by Putnam (1975) and Burge (1979, 1986)) since on the active version, were we (say, in some organismic twin) to retain the in-head structure but alter or remove the extended structure, the gross behaviour of the agent will change.

I do not propose to defend either active externalism or the extended mind thesis (EM) in the present treatment. But it has sometimes seemed that there must be a direct route from claims concerning the extended mind to claims concerning extended physical bases for the conscious mind (henceforth ECM). Thus Noë and Thompson write that:

Externalism about the vehicles of content (advocated in print by Hurley, 1998; Hurley and Noë, 2003; Rowlands, 2002, 2003; Clark and Chalmers, 1998; and Noë, 2004a,c), however does entail that neural systems are not sufficient for consciousness. (Noë and Thompson 2004a, p.94)

Notice that in the case of ECM, just as in the case of EM, what is at issue is the location of the material vehicles of certain mental or cognitive states. But whereas EM was concerned only with the vehicles of non-conscious mental states such as states of dispositional believing, ECM makes the even more striking claim that the local material vehicles of some of our conscious experiences might include more than the whirrings and grindings of the brain/CNS (Central Nervous System).

This focus on material vehicles (rather than on contents) also explains why there can be no simple inference from facts such as the presence of a feeling of

^{iv} Active externalism is in fact orthogonal to the more passive varieties. See Clark and Chalmers (1998).

pain ‘in the tooth’ to the conclusion that the apparatus underlying the conscious experience is (in the relevant sense) extended^v. For the question is not whether the states of affairs that the conscious thought *concerns* are extra-neural: they nearly always are. Nor is it a question about what individuates the thought as, for example, a thought about a tooth. What ECM – at least, the versions of ECM that I am setting out to target – sets out to challenge is a view about the location of the material underpinnings that enable the thought or experience to occur. It is a question about what might be dubbed the machinery of mind.

To take a simple example, perceived motion is not itself ‘in the head’: the motion is (usually at least) in the world. But we may still ask about the location of the material vehicles of motion perception. This is a question about the nature and location of the physical activity that *realizes* motion detection. It is this activity that is typically considered to occur ‘in the head’. Indeed, key elements of this machinery look to be even further localizable, to area MT, which can be selectively damaged so that the world is then perceived as a set of static moments (Marcar, Zihl, and Cowey, 1997). Where standard views in contemporary neuroscience depict all such key (local mechanistic) activity as taking place in the head/CNS, ECM intriguingly suggests that – for some forms of conscious experience at least – we should explore a larger material canvass, one that includes processing loops that reach out to embrace states or activity in the (non-neural) body and world.

It is this kind of claim that I want to examine. I hope to show that nothing in the arguments for EM should incline us to accept ECM (to accept an extended view of the mechanisms of the conscious mind or of the vehicles of conscious experience), that none of the other arguments offered by the proponents of the ‘processing loop’ versions of ECM fill in the gap, and that there is at least one good reason to think that no such argument (none predicated, that is, on the complexity or nature of the processing loops themselves) will be forthcoming.

It is worth stressing that the rejection of such ‘processing loop’ arguments leaves open the possibility of other (one might say, ‘more metaphysical’) arguments that might be thought to support something akin to ECM : for example, the arguments put forward by ‘naive realists’ (Martin, 2002, 2004) about perception. Those arguments offer a very different picture of the terrain, and might threaten (if successful) to blur the vehicle/content distinction itself,

^v Thanks to an anonymous referee for reminding me of the need to clarify these matters at the outset.

at least in the case of sensory perception^{vi}. Consideration of such arguments is beyond the scope of the present treatment.

I begin, then, by rehearsing (sections 3-5) the three main dynamic/enactive ('processing loop') arguments for ECM.

3. Sensorimotor Loops and Variable Neural Correlates

Noë (2004) suggests that in some way it is the shape of the whole sensorimotor loop that determines the character and content of perceptual (e.g. visual) experience. In developing this view he is pursuing the so-called 'enactive' approach according to which perceptual experience is *enacted* via an agent's skilled sensorimotor behaviour (Varela, Thompson and Rosch, 1991). According to the enactivist, 'Perception is not something that happens to us or in us, it is something we do' (Noë 2004, p. 1). The role of actual activity in these accounts is not, however, straightforward. For it is not activity itself, so much as the know-how that drives the activity, that ultimately plays the crucial role. Perceptual experience, so the story goes, gains its content and character courtesy of the exercise of sensorimotor know-how, that is courtesy of the active deployment of implicit knowledge of the relations between (typically) movement and sensory stimulation. Thus Noë writes that:

perceptual experience acquires content thanks to our possession of bodily skills. *What we perceive* is determined by *what we do* (or what we know how to do); it is determined by what we are *ready* to do ... we enact our perceptual experience: we act it out. (Noë 2004, p. 1)

As (partial) evidence for this view, Noë (2004) points to work involving Tactile Visual Sensory Substitution (TVSS) systems: work pioneered by Bach y Rita and colleagues (Bach y Rita, 1972). For a review, see Bach y Rita and Kercel (2003). In this well-known work, blind subjects are fitted with head- or eyeglass- mounted cameras that feed visual information to a small array of tactile stimulators mounted on the back (in the original work) or to a tiny electrically enabled tongue-pad (in the most recent versions). After a while, subjects report that the tactile sensations fade, to be replaced with quasi-visual ones able to support behaviours such as reflex ducking when a ball or other

^{vi} This is because they hold that "some of the objects of perception....are constituents of the experience" (Martin 2004, p. 39). Notice that this claim goes beyond the more standard 'direct realist' suggestion that when we perceive the world we do so without first *perceiving* something inner. The direct realist view is compatible with what Foster (2000) nicely dubs psychological (as opposed to perceptual) mediation. For some useful discussion see Millar (2007).

object is thrown towards the head. Importantly for the enactivist, learning about newly enabled sensorimotor loops (learning a set of ‘sensorimotor dependencies’, to use the terminology of Noë (2004), linking motor actions to new sensory stimulations) turns out to be crucial to success. It is only when a subject begins to learn about the ways actively moving the camera yields systematic changes in tactile input that the experience begins to seem ‘quasi-visual’ (e.g. she begins to experience looming). The agent’s experience then ceases to feel (only) like touch and starts to feel like vision. Such results lead Noë to a very strong claim. This is the claim that:

In general, what determines phenomenology is not neural activity set up by stimulation as such, but the way the neural activity is embedded in sensorimotor dynamic. (Noë 2004, p. 227)

This ‘embedding’ is said to have far-reaching consequences for the philosophical and scientific understanding of consciousness:

The enactive approach seeks to explain the quality of perceptual consciousness not as a neural function caused by and realized in the brain...but rather in terms of patterns and structures of skillful activity. On the enactive approach brain, body and world work together to make consciousness happen [...] Experience is not caused by and realized in the brain, although it depends causally on the brain. Experience is realized in the active life of the skillful animal. (Noë 2004, p. 227)

Following Noë, we may dub this the rejection of ‘neural sufficiency’. But such a rejection, I suggest, cannot be justified on the basis of the TVSS (and related) evidence. It cannot be thus justified because it depends on taking evidence for the role of whole sensorimotor loops in *training and tuning* the neural systems that support conscious perception for evidence of the ongoing role of such loops, or even just of implicit knowledge of such loops, in conscious perception itself. Nothing in the evidence makes a case for the latter claims. The evidence thus leaves open the possibility that embodied activity is just a causal precondition of the setting or re-setting of parameters in neural populations: parameters that, once set, suffice for the activation of those neural populations to bring about the experience in question (Block, 2005).

There is a subtler move possible hereabouts, and one that is pursued by Hurley and Noë (2003) and by Hurley (forthcoming). This move attempts to erode the importance of the training/post-training distinction by focusing not on the most local mechanism of occurrent experience itself, but rather upon the *best*

explanation of the quality and character of the experience. The suggestion is thus that we shift attention from what Hurley dubs the ‘sufficiency question’ (e.g. ‘what is the system within which a certain pattern of activity suffices for an experience of visual looming?’) to an explanatory one (in this case ‘why is this neural state the neural correlate of the experience of visual looming?’). Thus suppose we return to the example of TVSS. Following Hurley and Noë, we may suppose that after training and adaptation, at least some different neural regions are now implicated in the reconstituted ‘visual’ processing. Given such an outcome, we may reasonably ask what it is that standard vision and TVSS-supported vision have in common? But the answer, they suggest, is not apparent from the neural data alone. Instead, what they have in common is their ability to support what Hurley (forthcoming) describes as ‘a characteristic extended dynamic’: a distributed process involving brain, body and the active probing of the world. It is this sameness of extended dynamic pattern, Hurley argues, that best explains the sameness of experiential quality. So even if activity in the neural stuff alone suffices (after training) for the experience, our explanations of the visual-qualitative nature of the experience need to look further afield.

But this version of the argument, though it avoids the previous worry, then fails to provide any support for ECM. To see this, consider that even standard teleosemantic forms of representationalism^{vii} (which identify contents, even when neurally encoded, by what they are about) could avail themselves of this argument to place various neural states into a content-based equivalence class. Such a result is interesting, but falls far short of undermining standard internalist views about the local (neural) vehicles of content. Or suppose what matters is the achieved functional poise of a representational state? Then the very same content and the very same poise might be neurally supported in a variety of ways (perhaps, as in TVSS, reflecting their recruitment via different gross input channels). But once again, there would be no threat posed to standard views about the location or nature of the machinery of mind.

But perhaps there is something special about the specific application of the variable neural correlates argument to the case of perceptual experience? Hurley and Noë (2003) claim that in this arena their view has a clear advantage over representational ones, since:

when it is brought to our attention that certain sensorimotor contingencies are characteristic of vision ... it [becomes] intelligible why

^{vii} Classic examples include Dretske (1981) and Millikan (1984).

it is like seeing rather than hearing to perceive in a way governed by the sensorimotor contingencies characteristic of vision rather than [those of] audition. (Hurley and Noë 2003, p. 146)

The idea here is that learning that some neural activity pattern P in area Q correlates with visual experience leaves us wondering why: why does all that support visual (rather than e.g. tactile) experience. By contrast, it is argued:

when the sensorimotor pattern characteristic of vision is explained, we have an ‘aha’ reaction; we see through the dynamic pattern of sensorimotor contingency to what vision in particular is like. (Hurley and Noë 2003, p. 160)

But this threatens to underestimate the resources of the standard approach (and, correlatively, to overestimate those of the proposed alternative). For a good neurocentric account need not just nominate an area or activity pattern as the neural correlate of an experience or type of experience and leave it at that. It may also make clear why, given that (for example) such and such information is now functionally poised in such and such a way, the agent will tend to say and do the very things they tend to say and do: the very things characteristic of, for example, seeing not hearing. Such behaviours may include reporting looming, ducking on receipt of incoming-baseball-specifying visual information, and so on. Such a result would merit at least a mild ‘aha’! It might be thought nonetheless to fall short (though see Dennett (1991) for a famous rejection of this alleged shortfall) of explaining why visual experience has the specific qualitative character it has. But the enactive/variable neural correlates story is equally silent on this very point, as Hurley and Noë (2003) candidly admit.

I conclude that there is nothing in the simple commitment to enactivism (or in its more epistemologically spun cousin, the Variable Neural Correlates argument) that supports ECM.

4. Virtual Representations

A second kind of argument meant to favour ECM takes as its starting point some well-established facts (for a handy review, see Stafford and Webb (2005, p. 38-42)) concerning the limits of here-and-now perceptual uptake. For example, the visual pick-up of colour information does not extend to the edge of our visual contact with world, since the colour sensitive cells are almost all in the densely photoreceptor populated central high resolution fovea (although

brightness can be detected all the way out, and sometimes we get colour cues from that). In addition, visual sensitivity to detail is high only in that small foveal window. Yet, as Dennett (1991) and others note, we do not experience the world as only detailed and coloured in the middle! Why not?

One possible answer is that the scene appears coloured and detailed ‘all the way out’ because we mistake easy accessibility for actual encoding. Thus we implicitly know that we can retrieve colour information and finer and finer detail at will, just by moving our heads and eyes so as to scan the scene via a series of rapid eye movements known as saccades. The most common way to unpack this suggestion is to suggest that the visual experience of richness of detail etc is illusory: we *think* we experience rich detail (etc) but we do not (Dennett, 1969, 1991; Ballard, 1991; O’Regan, 1992; Churchland et al, 1994). Instead, we simply represent to ourselves the fact that the scene is full of accessible detail.

But an alternative (and for present purposes more interesting) account suggests that the experienced richness is not an illusion at all. Rather, the correct lesson (it is argued) is that our perceptual experience is determined not just by the current neural encodings and activity but by the *combination* of those encodings and activity, our own capacities for saccadic action, and the actual detail of the external scene (Noë, 2002, 2004, 2007, 2008). The colour, detail etc, is (as Noë puts it) ‘virtually represented’:

the world is present in experience virtually, *the way information from a remote server is present on your desktop*. The world is present virtually thanks to the way we are bound to it, in bodies with the right sort of networked connections. The flick of the eye, the turn of the head, the movement of the body, brings us the detail we need as we need it. *The world is present virtually thanks to our online, dynamic access to it*. (Noë 2007, ms p. 15)

The idea once again is that visual experience is enacted, insofar as it is partially constituted by our actions and by the world we act in. The best explanation of our experiences of detail (etc), so the argument goes, is that such experiences, just as ECM requires, are not determined by inner neural activity alone.

There is, however, a problem. For as Noë himself notes the remote server analogy is not necessarily a good tool for securing this conclusion. It seems to suggest that ‘experience has this content only as a potentiality’ (Noë 2007, p. 15). But mere potentialities of experience are surely not what is at issue between the friends and foes of ECM. The target is experience itself, and the

question is, what are the local physical goings-on that determine the nature of that (actual, not potential) experience? Here, Noë makes an interesting but ultimately rather puzzling suggestion. He suggests that (unlike in the computer case) ‘you cannot factor experience into an occurrent and a merely potential part’ (Noë 2007, p. 16) since:

Pick any candidate for the occurrent factor. Now consider it. It is structured too; it has hidden facets or aspects too. It is present only in potential. (Noë 2007, p. 16)

Adding that:

... this is perhaps the most important idea in this paper, experiential presence is virtual all the way in. This is an important disanalogy with the computer case. (Noë 2007, p. 16)

Thus it is agreed on all sides that we seem to see colour all the way out, but that our here-and-now pick up is more limited. But now, Noë (2007) suggests, try and attend to (for example) just that part of the colourful shirt you have in true (occurrent, here-and-now) foveal view. It will turn out that that part of the visual scene too has structure, and hence (according to Noë) that it seems as it does only because you can attend to *its* parts as needed too. And so on all the way in, whatever location and feature you choose. Experience, Noë claims, is thus ‘virtual all the way in’. So unlike in the computer case, there is no contrast between what is truly experienced (what is already locally encoded, as it were) and what is experienced in virtue of potentialities (what is available by online dynamic access).

I confess to being unsure how to understand this argument. The original ‘mere potentialities’ version of the virtual representation story strikes me as generally plausible (Clark, 2002), but it does not (as Noë seems to admit) actually support ECM. The trickier ‘virtual all the way in’ version attempts to avoid this shortfall, but does so (it seems to me) at the cost of considerable obscurity. Is the claim just that we can always attend to, and hence bring into experience, more detail? If so, it is not at all clear why that supports anything like ECM. Is it perhaps that what fixes any experience is not a snapshot moment of neural activity but a process extended in time? That sounds plausible (and we shall return to it later). But then – at least at first pass - all that seems to matter is that the neural activity evolve over time in such-and-such a way. In which case the machinery of conscious experience is all in the head, with the proviso that the head persists over time and that the neural encoding evolve in a certain

way. This too falls far short of establishing ECM. As it stands, the ‘virtual all the way in’ version of the argument is either flawed or simply too opaque to carry the weight required.

I conclude that the appeal to virtual representing fails as an argument for ECM.

5. Dynamic Entanglement

This brings us to the last (and most promising) of the current arguments for ECM: the argument from dynamic entanglement. The starting point for this argument is the idea (increasingly influential in the sciences of mind) that we should, in many cases, resist the temptation to think in terms of a simple linear flow in which the senses deliver input which is progressively processed and refined until an output (usually a motor action) is selected, and the process repeats. This picture (which Hurley (1998) dubs the Input-Output Picture) has been challenged on many grounds^{viii} but the key observation is that motor processing and perceptual uptake each unfold courtesy of a mass of ongoing looping interactions in which recurrent neural circuitry and bodily action combine so as the active agent structures the information flow in ways apt to the task (for more on the self-structuring of information flows, see Lungarella and Sporns (2005) and Clark (2008)).

An alternative to the simple input-output model thus stresses the looping dynamics of the processing, describing cognitive mechanisms in which ‘information flows back as it flows up, and it flows more or less continuously’ (Hardcastle 1998, p. 341). The physical vehicles are, on these accounts, sometimes extended dynamic loops connecting ‘higher’ to ‘lower’ brain areas, and encompassing both ‘cognitive’ and ‘motor’ systems, (see also Clark (1997)). Clark (1999) dubs this alternative vision ‘Escher Spaghetti’ where this names a seething mass comprising not just multiple criss-crossing strands (ordinary spaghetti), but strands whose ends feed back into their own (and others) beginnings, making ‘input’ and ‘output’, and ‘early’ and ‘late’ into imprecise and misleading visions of complex recurrent and reentrant dynamics.

To turn these kinds of observations about complex neural dynamics and active perception into an argument for ECM, however, we need to add some further ingredients. In particular, we need to combine the picture of complex looping processes with the claim that the conscious-experience supporting loops (providing the non-instrumental minimal necessary supporting structure for at

^{viii} See the discussion in Hardcastle (1998). For a review, see Clark (2001) chapter 5.

least some forms of conscious experience) are those running (not just within the brain but) through brain, extra-neural body, and perhaps even world. Versions of just such an argument appear in Hurley (1998), Thompson and Varela (2001), Noë and Thompson (2004a,b), Noë (2008), and Cosmelli and Thompson (forthcoming). Thus Thompson and Varela write that:

[W]e conjecture that consciousness depends crucially on the manner in which brain dynamics are embedded in the somatic and environmental context of the animal's life, and therefore that there may be no such thing as a minimal internal neural correlate whose intrinsic properties are sufficient to produce conscious experience. (Thompson and Varela 2001, p. 425)

In just this vein Cosmelli and Thompson (forthcoming) argue that the contributions of the (non-neural) body are so important and complexly intertwined with the neural processing itself that we cannot simply 'carve off' the neural elements from the rest. The root cause of this is supposed to involve a certain kind of dynamic complexity sometimes known (Clark, 1997, 2002; Wheeler, 2005) as continuous reciprocal causation. This kind of complexity is found in 'dense nonlinear systems where all state variables interact with each other, any change in an individual variable becomes inseparable from the state of the entire system' Cosmelli and Thompson (forthcoming). Cosmelli and Thompson offer lots of neat examples of this kind of complexity, but the central idea emerges most clearly in the following (rather long) passage:

[the examples] are intended to stress the immense complexity of the neural and extraneural interactions that ultimately determine brain activity in the living organism. The list of functional systems dependent on brain-body coupling to provide the organism with coherent perception of the world also includes the entire interoceptive, autonomic system ... vestibular-autonomic regulation ... balance and somatic graviception relying on hydrostatic properties of blood pressure and inertial mass of abdominal viscera ... , as well interaction between the senses occurring at both central and peripheral levels. (Cosmelli and Thompson, forthcoming, **ms** p. 12)

Experiences of emotion, to take just one possible example, are sometimes said to depend on highly complex, temporally extended processes looping between

brain and the extra-neural body (Varela, 1999; Thompson and Varela, 2001)^{ix}. But once we accept that kind of model for the experience of emotion, it seems a small step to accepting it for the experience of seeing, and thus extending the relevant loops out into the wider world. Thus work in ‘active’ and ‘animate’ vision (Ballard et al, 1997) stresses the role of bodily acts (such as head motions and saccades) and the persisting real-world scene in visually-based problem solving. Does visual experience itself, as ECM suggests, depend non-instrumentally on such body- and world- involving processing loops?

On the face of it, the answer is ‘no’, or at least, ‘not proven’. For intuitively, visual experience may at each moment depend solely on the complex brain activity caused (instrumentally) by the ongoing engagements with body and world. The same might be said for the case of emotion, if various body-involving loops matter only insofar as they ‘report back’ to the brain and CNS. The ‘dynamic entanglement’ arguments for ECM do not, however, stop there. For there is one more move to be considered, and it is (it seems to me) the crucially important one. It concerns the vexed question of temporal spread.

Suppose that we reject what Noë (2004, p. 35-39) calls the ‘snapshot picture’ of visual experience, as fully determined by brain states at some moment in time, and instead suggest that it is *processes* that suffice for conscious experience and that these processes essentially evolve in time? In that case, there may seem room for certain forms of looping (body and world involving) engagement to play a constitutive rather than an instrumental role. According to such a view:

^{ix} But notice that here too (recall section 2 above) we need to be wary of the exact claim that is at issue. For while it may be compelling to conceive of certain emotions as essentially priming bodily responses to threats, dangers, risks etc (and of the experience of touch, to take another such example, as in some way essentially bodily) it is by no means evident that the most local material vehicles of emotional experience (or touch) involve processing loops that extend beyond the brain/CNS. Thus Jesse Prinz, in a recent treatment entitled ‘Is Consciousness Embodied’ writes that:

The claim that consciousness extends into the body is only marginally more plausible than the claim that consciousness leaks out into the world. We have never found any cells outside the brain that are candidates as correlates for experience. Such cells would have to co-vary with conscious states in content and time course. Every component of the body that we can experience is represented in the brain, and when the corresponding brain areas are damaged experience is lost. Conversely, bodily experience can continue after the body is damaged, as in the case of phantom limb pain. There is, in short, little reason to think the correlates of experience extend beyond the cranium. (Prinz 2008, p. 425)

What I am calling the ‘processing-loop versions’ of ECM aim to provide just such reasons. It is these arguments that are at issue in the present treatment.

experience ... is a temporally extended skill-based activity, comparable to the playing of a game, or a dance. ... There's no such thing as my experience at an instant in time. At a given time I can report on my experience. But what I do then is not record my state at that instant, but report on the evolution of my engagement with my surroundings at that instant. Experience is a temporally extended activity essentially. (Noë 2008, p. 460)

The upshot of all this, according to Noë, is nothing less than ECM itself, viz, that 'it is not the brain alone ... that suffices for experience' (Noë 2006, p. 420).

Does ECM really follow? The best test of this view is still the 'twin test' mentioned in connection with the arguments for EM and active externalism rehearsed in section 1 above. Thus take a here-and-now neural duplicate of you and ask; Would the duplicate, instantaneously ushered into being by, let us imagine, some quantum accident, share your experiential state? The answer, if the 'temporal spread' considerations are correct, must be 'no'. Indeed, Noë suggests that a negative answer to this question follows pretty much directly from an enactivist conception of perception:

I have urged that experience is a temporally extended phenomenon; it is an activity of skillful probing. If this is right, then *a neural duplicate of me now, at a moment in time, won't, by dint of being my duplicate now, have any experience at all*. If the duplicate does have experience, it will be *thanks to its dynamic, temporally extended interaction with the environment*. But then again we must note that there is little reason to think that its experience would or could be like mine unless its environment were also like mine. (Noë 2006, p. 420, my emphasis)

Suppose we agree (as seems independently plausible) that there will be no experience without some temporal evolution of brain states. We thus agree that for information to become conscious, some amount of time needs to pass, so that normally there is no way the brain can 'in an instant' reach the kind of state that supports conscious experience. This fact alone does not mandate ECM. For it might still be the case that what normally takes time is for activation to build up (perhaps courtesy of recurrent neural circuitry (Lamme, 2006; Block, 2007)) until some kind of threshold is passed. If that were so, whatever builds up during the normal time frame might still be extraordinarily ushered into being by our quantum accident. This kind of role for temporal evolution is thus fully compatible with both a neuro-internalist and essentially snapshot conception of the physical underpinnings of conscious experience. We can

even relax, if we wish, the snapshot element here while retaining the internalism. Thus imagine that even once the right kind of activation state is reached, the state still needs to persist for a while for a conscious percept (say) to be experienced. Even so, that still gives us no reason (yet) to look outside the brain for the sufficient material substrate of the experience.

But perhaps the point is that the neural states, to support a given experience, need not just to persist but to evolve, over time, in some specific way? Here too it is tempting to defend the internalist intuition, by arguing that whatever this signature evolution may be, it is surely an evolution of neural states, and so all we need to do is magnify our quantum accident so as to bring about that very same state evolution to ensure (purely internally) that the experience occurs. Once again, the minimal substrate of the conscious experience looks to be attainable regardless of the states of the extra-neural body or extra-bodily world.

It is at this point that things become murky, and interesting. For, as Hurley (1998) points out, it is open to the naturalistically inclined philosopher to reject the thought experiment as an indicator of the minimal circuitry of experience. Perhaps some specific experiences (not all experiences, let us suppose, but some) require a kind of ‘signature’ temporal evolution of neural states that simply cannot (in the natural order) occur in the absence of the right extra-neural scaffolding. Thus consider, as a rough parallel, the case of a large orchestra that cannot play *quite like that* without the conductor (for ECM the ‘conductor’ is, of course, the loops via body and world). In this vein Noë suggests that:

... perhaps the only way – or the only biologically possible way – to produce just the flavor sensation one enjoys when one sips a wine is by rolling a liquid across one's tongue. In that case, the liquid, the tongue, and the rolling action would be part of the physical substrate for the experience's occurrence. (Noë 2004, p. 220)

Similarly, perhaps the rich visual scene can never look *quite like that* unless you are really acting and behaving in the world. Would this be constitutive or merely instrumental dependence? I do not think the answer is clear-cut. But at this point we have at least isolated what seems to be the only plausible argument from dynamic entanglement to ECM. It may be represented as a kind of equation:

(DEUTS: Dynamic Entanglement plus Unique Temporal Signature Argument)

Deep Dynamic Entanglement + Unique Temporal Signature = ECM

The ‘dynamic entanglement plus unique temporal signature’ argument (henceforth DEUTS) captures most of what Cosmelli and Thompson (in press) say about profound bodily involvement in the construction of experience, and of what Noë (2004, 2007, 2008) says about worldly involvement. Should it convince us of ECM?

6. Some Common Responses

One common but unconvincing, response to these arguments for ECM involves the appeal to the ‘brain in a vat’^x. Take whatever neural system you favor (the neural circuitry of the brain perhaps, or of the brain plus CNS). Keep it alive in a bath of nutrients, and provide it with inputs that faithfully simulate the energetic patterns that impinge on normal embodied, situated brains, and monitor its outputs so as to alter the contents of the simulation accordingly (allowing it to act upon and alter states of the simulated environment). For all we know, so the argument goes, we might be such brains, unknowingly envatted, and apparently acting on the world around us, apparently deploying yellow sticky notes, rolling wines across our tongue, using notebooks, etc. Does this not show that there can be no constitutive involvement for physical machinery beyond the bounds of the brain/CNS?

The problem with this argument is that it proves too much. To see this, imagine a variant case (Clark, 2008) in which a partially lesioned brain is envatted, and in which the helpful scientists use a cleverly (deep-neurally) hooked-up vat to fill in the usual contribution from neural motion area MT. The envatted agent now experiences motion. Presumably, this having of motion experience ‘by the envatted brain’ would not then show that activity in MT is not part of the normal physical substrate of motion experience! The mistake, then, is to infer that the sufficient mechanism is the biological stuff

^x This thought experiment is discussed in, for example, Smith (1984) and Brueckner (1986). An especially entertaining version appears in Dennett (1981). For some contemporary updates see Chalmers (2005) and Clark (2005)

alone, just because the biological stuff, in the special vat-context, helps support thinking and experience. At the limit of this thought experiment we have the single neuron in a dizzyingly complex vat (Wilson and Clark, 2008). We would not conclude that experience and thought constitutively depend only on the activity of that single neuron!

The trouble, in short, is that the intelligent vat (as deployed against ECM) is here asked to do all the complex work of body, action, and world (Hurley (forthcoming); Noë, 2004). This leads to a kind of dilemma. If the vat *does not* fill in everything the world provides, the experiment is unfair. If it *does*, it cannot prove anything, as the filled in contributions might (as in the case of the vat-repair of MT) still be essential for that very experience. The same considerations apply to all cases where a full simulation recreates every effect of bodily motion, somatic signaling etc, thus building in every functional effect of the world and body anyway (Noë, 2007; Cosmelli and Thompson, forthcoming). The brain-in-a-vat considerations are thus unable to advance the argument.

Another common, but equally inadequate, response to arguments for ECM is to point out that we can sometimes have experiences without the usual involvement of the body and world. This is strikingly so in the case of dreams, or when the probing neurosurgeon stimulates bits of tissue causing us to hear the opening bars of a symphony, smell the jasmine and *madre de la noche* from an old Andalucian holiday, etc. Thus (so the argument goes) since *some* experience does not need the active body and the wider world, maybe *no* experience does? Clearly, this is an unwarranted conclusion. To see this, we need only note that just because some of my experience does not involve e.g. auditory cortex, that does not mean that none of it ever does! Worse still, the cases where we do have the active body and wider world ‘in the loop’ way outnumber the others.

Nor does it help simply to discover neural correlates of conscious experience, be they ever so fine-grained. Thus suppose we find all manner of interesting neural correlates for specific kinds of conscious experience. Suppose even that these correlates can be tracked as they evolve moment-by-moment sensitively linking brain activity to the agent’s unfolding experience. For example, Schyns et al (2007) use EEG signals to track (moment-by-moment) the key state transitions in the brain as different bodies of information (those specifying ‘nuns’ versus ‘Voltaire’ in the two visual interpretations of the famous ambiguous picture) are processed, one step at a time, in the brain. The upshot

is a compelling picture of the evolution and significance of the crucial neural processes over time.

What follows regarding ECM? Not much, it seems. For all this neural activity delicately unfolds in the presence of the ambiguous picture itself. So how do we know that the picture and our saccades around it do not form part of the minimal sufficient machinery for that very experience? Those who reject ECM assert that the picture and saccades merely drive the brain through a sequence of states, but that the brain processes themselves provide the sufficient machinery of the experience. By contrast, those who embrace ECM claim that there is no good reason, independent of our neurocentric prejudices, to make the cut between essential ongoing causal ‘drive’ (recall Noë’s case of the taste of the wine and the rolling of the liquid over the tongue) and experience-sustaining machinery at that very point. Since we cannot (given our earlier argument) invoke the experience of the brain-in-a-vat to support the standard causal/constitutive cut, how are we to resolve this?

We do not make progress by simply asserting that the cut be made at the brain/body (or even organism/world) boundary. Nor do we make progress by simply conflating, as proponents of ECM sometimes do, the notions of causal drive and essential machinery themselves. Noë sometimes looks to be guilty of this error, in passages such as these:

According to what Clark and Chalmers [1998] call active externalism, the environment can *drive and so partially constitute* cognitive processes. (Noë 2006, p. 411, my emphasis)

if ever there was a plausible candidate for a psychological state that is *driven and so partially constituted by the environment*, it is perceptual consciousness. (Noë 2008, p. 460, my emphasis)

But this misrepresents, or at any rate uncomfortably oversimplifies, the Clark and Chalmers argument for EM. For Clark and Chalmers nowhere suggest that *just because* X drives Y, X becomes partially constitutive of Y. That would indeed be to make a ‘causal-constitutive’ or ‘coupling-constitutive’ error (Adams and Aizawa, 2008; Block, 2005). Just because my TV picture is sensitively driven by the incoming signal, that does not make the transmitter part of the minimal sufficient physical substrate for that very picture (not even as the picture evolves over time). Rather, in the Clark and Chalmers argument for EM, everything depends on the *precise way* X and Y together function to control behavior (the way, we argued, distinctive of dispositional believing). We

should thus reject any arguments (for EM or ECM) that merely conflate causal drive and essential phenomenon-producing machinery^{xi}. For they secure ECM only at the cost of a radically unconvincing general principle (that to sensitively drive some X is to be part of the machinery in virtue of which X obtains).

At this point in the dialectic, a deep stalemate beckons. Just because OUTER drives BRAIN, and BRAIN depends (let us assume) on OUTER to step through the signature sequence of states that support some specific experience, that does not yet show that OUTER is part of the minimal machinery of experience. But nor can we simply claim the opposite (given the failure of the standard vat-style thought experiments) except as an expression of our pre-existing prejudices.

7. Keeping Joy In Its Place

There is one final set of considerations, however, that may yet begin to untangle these argumentative kites. The considerations concern timing (again) and bandwidth. Thus Chalmers (2008) suggests that arguments for EM may fail to generalize to ECM and that:

Perhaps part of the reason is that the physical basis of consciousness requires direct access to information on an extremely high-bandwidth [...] our low-bandwidth conscious connection to the environment seems to have the wrong form as it stands. (Chalmers 2008, p. xii-xii)

Chalmers (2008) does not develop this suggestion, but the direction seems promising. Perhaps conscious awareness is special among cognitive functions insofar as it requires (in us humans at least) certain information-accessing and information-integrating operations whose temporal scale makes neural (brain/CNS) processes (just as a matter of contingent fact, in us humans) the only adequate ‘vehicle’. (Note that the Clark/Chalmers case for EM, by contrast, targets only non-conscious mental states and processes, where long-term informational poise (rather than online informational access and integration) seems to be what counts).

As a conjecture about the physical roots of conscious experience, this view has some plausibility. Thus the philosopher and neuroscientist Chris Eliasmith (2008) suggests that the dynamics internal to the brain are ‘qualitatively

^{xi} Adams and Aizawa suggest that the Clark and Chalmers arguments for EM are equally guilty of this conflation. I respond in more detail to this charge elsewhere (Clark, 2007, 2008) and will not further repeat those arguments here.

different' from those that span the brain-world boundary and that the key differences involve speed and bandwidth:

The most obvious differences [between neural and super-neural dynamics] are the *speed of information flow (i.e., bandwidth)*, and the *degree and kind of coupling*. Because bodies have mass, they tend to *slow down the transfer of information to the world from the brain* (i.e. they effectively act as a *low-pass filter*). However, *no such impediment to information flow* exists between brain areas. This results in a huge difference between the kinds of coupling that can be supported between brain subsystems and between the brain and the external environment. (Eliasmith 2008, p. 150, my emphasis)

When does such a difference make a difference? Not, we can reasonably assume, in the case of non-occurrent states such as dispositional believings^{xii}. But it is plausible that speed (or fine temporal issues more generally) makes a crucial difference in the moment-by-moment construction of conscious experience itself. Thus suppose conscious experience requires cortical operations that involve extremely precise temporal resolutions, such as the synchronous activation of distinct neural populations where the required synchrony demands millisecond precision (for some of the evidence for this conjecture, see the review in Singer (2003)^{xiii})? This might be so, for example, if the brain uses fine temporal synchronization and fast signal processing to bind together sensibly coherent bodies of information. Or alternatively, as suggested by Thiele and Stoner (2003), and by Lamme and Spekreijse (1998), it may be that synchrony correlates with attention, and that attentional modulation is what allows information to pass from perceptual buffers to working memory in the way that gives rise to conscious experience^{xiv}.

In all these cases, the fans of ECM will suggest that we are still studying only the *neural component* of the true substrate of experience. But perhaps we can now see a principled reason to be skeptical. In such cases the external environment may well *matter* insofar as it causally drives the neural systems, but the key effects that enable and explain the quality of the felt experience may be occurring at time-scales that are only possible within the neural apparatus itself.

^{xii} Eliasmith (2008) presents the temporal considerations as a general reason to be skeptical of claims concerning true cognitive extension. But the empirical findings that best support the timing argument concern only the construction of conscious experience, and depend essentially upon the need for fine time-scale synchronies as a means of binding together bodies of neurally represented information.

^{xiii} Thus we read that 'Cortical neurons can engage in oscillatory firing patterns...and synchronize their responses with millisecond precision over surprisingly large distances'. (Singer, 2003)

^{xiv} Thanks to Jesse Prinz for alerting me to this possibility.

If this were so then everything that involves subsequent motor responses or bodily actions (including active saccades around the scene) will be ‘screened off’ (by the bodily ‘low-pass filter’) from the neural/CNS mechanisms that actually produce the conscious experience. This is the key effect that may defeat the most promising argument (the DEUTS, ‘dynamic entanglement plus unique temporal signature’ argument) for ECM. It is worth pausing, then, to clarify the shape of the worry.

A low-pass filter is any physical medium that allows low frequency signals through while reducing or blocking higher frequency signals. The walls of a room act as a low pass filter for sound. That is why you hear the low-frequency bass more than the high-frequency treble from the hi-fi when you are in the next room. Eliasmith’s interesting suggestion is that the extra-neural body, implicated in all cases of active vision and motor loops, acts as a kind of low pass filter for signals coming from the environment. What this means in practice is that for phenomena that depend on, for example, the very fast temporal binding or processing of signals, the only locus in which such operations can (as a matter of fact) occur lies within the brain/CNS. The muscles, for example, would act as a low pass filter, and activity there would thus fall outside the effective system within which signal binding or processing on the right timescale can occur. Muscular goings-on, if this is correct, could indeed be a source of inputs to the system that generates conscious experience (as they must be, given that we can experience muscular action) but they will not form part of the system upon which the experiences most locally depend. Muscular activity (like environmental signals in general) would thus be fit to play a causal rather than a constitutive role in the construction of experience^{xv}.

This is actually quite an intuitive result. Thus Adams and Aizawa (2008) write, concerning arguments for the extended mind in general, that:

The orthodox might listen to an argument for the view that cognition extends beyond the neurons of the brain into the spinal cord and

^{xv} The same considerations apply in the case of another worry usefully raised by an anonymous referee. The worry is that ‘the way we hear ourselves speak contributes constitutively to the experience of speaking’ since it involves ‘both a sense of the movement of one’s vocal organs and of the sounds one thereby makes’. Once again, the key is to separate two forms of non-instrumental involvement for these bodily aspects. For it may well be that the *sense* of both movement and sound form essential aspects of the experience of hearing ourselves speak. But the question before us is whether the most local processing activity that *suffices* for that sense is itself spread out, or involves only neural encodings and operations. And the empirical story I have sketched offers a clear answer: the local operations that suffice for it seeming to the agent that she is moving her vocal organs and hearing her own voice are all staunchly contained within the neural apparatus, even though the signals that drive that apparatus are (unless she is hallucinating) rooted in peripheral bodily events. (See also note 7 above)

sensory nerves, but muscular cognition is beyond the pale. The standard assumption is that the kinds of cognitive information processing that take place in nerves are dramatically unlike those that occur in muscles. (Adams and Aizawa 2008, p. 19)

The problem with this as an argument against EM is that we need to know what kind of difference is in question, and just why it *makes* a difference for the obtaining of some specific mental or cognitive state. It is not enough (as nicely argued by Wheeler (in press)) merely to note some difference or other, on pain of rapidly begging the question against the very possibility of non-neural components of cognitive processes^{xvi}. But in the special case of arguments for ECM, we can begin to discharge this obligation. For if indeed the physical machinery of conscious experience requires fast time-scale operations and processing, and the non-neural body acts as a low pass filter preventing external (and internal muscular etc) signals from directly entering into such operations and processing, then such signals are fit to play only a causal role, driving the neural systems within which the right kinds of fast binding and processing can occur. In such cases one might have all manner of complex couplings without thereby producing an extended material base for conscious experience. Contrast the case (discussed at length in Clark (2007, 2008)) of the possible role of gesture in the process of reasoning. There seems no reason why slow time-scale gestural events should not productively interact with faster time-scale neural ones so as to yield a special kind of coupled gestural-neural unfolding that is *itself* the distinctive physical engine of a certain kind of problem-solving. But within this coupled unfolding, the streaming contents of conscious experience would all depend constitutively only upon the neural processing itself. The account on offer thus enables us to embrace the kinds of claim made by Noë and others to the effect that certain experiences may only come about due to the neural systems being driven, in some distinctive way, by external signals. But it does so without being forced to the conclusion that such external sources comprise part of the most local machinery that generates the conscious experience itself. The account thus offers a principled reason for making the causal/constitutive cut, in the special case of conscious experience, in an orthodox, non-extended, kind of way^{xvii}.

8. Conclusions: Walking The Line

^{xvi} Adams and Aizawa do try to provide such reasons, though (again, see Clark (2007,2008) for discussion) the reasons are unconvincing in the case of non-conscious mental states.

^{xvii} This result obtains only for standard human agents circa 2009. Future direct Brain-Machine Interfaces (Serruya et al, 2002) may change all this. Properly ‘jacked-in’ via some future fast, broad-bandwidth interface, we might yet expand the physical substrate of conscious experience itself.

Might the most local machinery whose activity is sufficient for conscious experience include more than the brain/CNS? We have explored a variety of ‘processing loop’ based arguments meant to suggest this conclusion, and found them all wanting. The most promising such argument, DEUTS (the argument from dynamic entanglement plus unique temporal signature) fails in an especially revealing way. For if the empirical considerations advanced in the previous section are correct, then we can at least in principle walk the difficult line that the best arguments for ECM seek to challenge. For we can allow that at least some conscious experiences may have the precise qualitative nature they do only when the brain/CNS is being sensitively driven by the body/environment in a specific way, while still maintaining that the involvement of the body/environment is here merely causal rather than genuinely constitutive. We can do this in virtue of a specific but popular class of theories concerning the way the human brain constructs ongoing conscious experience. This is the class of theories that require fine-grained processes of temporal coordination to bind together, or otherwise process, sensibly coherent bodies of represented information. Should such theories (or indeed any others that have the same kinds of temporal implications) prove correct, DEUTS would fail. DEUTS, however, was the only promising argument for ECM. I conclude that the case for ECM is at best unproven and that the machinery of conscious experience is (probably) all in the head/CNS^{xviii}.

Department of Philosophy,
School of Philosophy, Psychology, and Language Sciences
University of Edinburgh,
EH8, 9AD, UK

ANDY CLARK

^{xviii} This paper (though it takes an opposing view) was inspired in large part by the work of Susan Hurley. She is greatly missed. Thanks too to Rob McIntosh, Alva Noë, Julian Kiverstein, David Chalmers, Michael Wheeler, the Edinburgh CONTACT group, and all the members of the Edinburgh Philosophy, Psychology and Informatics Reading Group for stimulating discussions. Thanks also to an anonymous referee for this journal, and to Keith Allen, Dave Ward, Jesse Prinz, and Georg Theiner for their comments on an earlier version. This paper was prepared and completed thanks to support from the AHRC, under the ESF Eurocores CNCC scheme, for the CONTACT (Consciousness in Interaction) project, AH/E511139/1.

References

- Adams, F. and Aizawa, A. 2008: *The Bounds of Cognition*. Blackwell: Oxford.
- Bach y Rita, P. 1972: *Brain Mechanisms in Sensory Substitution*. Academic Press: NY.
- Bach y Rita, P. and Kercel, S.W. 2003: 'Sensory Substitution and the Human-Machine Interface'. *Trends in Cognitive Sciences*, 7:12, pp. 541-546.
- Ballard, D. 1991: 'Animate Vision'. *Artificial Intelligence*, 48, pp. 57-86.
- Ballard, D., Hayhoe, M., Pook, P., & Rao, R. 1997: 'Deictic codes for the embodiment of cognition'. *Behavioral and Brain Sciences*, 20, pp. 723-767.
- Block, N. 2005: 'Review of Alva Noë, Action in Perception'. *Journal of Philosophy*, CII, 5, pp. 259-272.
- Block, N. 2007: 'Consciousness, accessibility, and the mesh between psychology and neuroscience'. *Behavioral and Brain Sciences*, 30: 5-6, pp. 481-499.
- Brueckner, A. 1986: 'Brains in a Vat'. *Journal of Philosophy*, 83, pp. 148-67.
- Burge, T. 1979: 'Individualism and the Mental', in P. French, T. Uehling Jr., and H. Wettstein (eds.), *Midwest Studies in Philosophy*, vol.4, Metaphysics. Minneapolis, MN: University of Minnesota Press.
- Burge, T. 1986: 'Individualism and Psychology'. *Philosophical Review*, 95: 3-45.
- Chalmers, D. 1995: *The Conscious Mind*. Oxford University Press: NY.
- Chalmers, D. 2005: 'The Matrix as Metaphysics', in C. Grau (ed) *Philosophers Explore The Matrix*. Oxford University Press, 2005. Also appears on the official Warner Bros Matrix website at: http://whatisthematrix.warnerbros.com/rl_cmp/phi.html
- Chalmers, D. 2008: Forward to Clark, A. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press: NY, pp. ix-xvi.

- Churchland, P., Ramachandran, V., & Sejnowski, T. 1994: 'A Critique of Pure Vision', in C. Koch & J. Davis (eds) *Large-Scale Neuronal Theories of the Brain*. Cambridge, MA: MIT Press, pp. 23-61.
- Clark, A. 1997: 'The Dynamical Challenge'. *Cognitive Science*, 21:4, pp. 461-481.
- Clark, A. 1999: 'Visual Awareness and Visuomotor Action'. *Journal Of Consciousness Studies*, 6: 11-12, pp. 1-18.
- Clark, A. 2001: *Mindware: An Introduction to the Philosophy of Cognitive Science*. Oxford University Press: NY.
- Clark, A. 2002: 'Is Seeing All It Seems? Action, Reason and the Grand Illusion'. *Journal Of Consciousness Studies*, 9: 5/6.
- Clark, A. 2005: 'The Twisted Matrix: Dream, Simulation or Hybrid?', in C. Grau (ed) *Philosophers Explore The Matrix*. Oxford University Press 2005. Also appears on the official Warner Bros Matrix website at: http://whatisthematrix.warnerbros.com/rl_cmp/phi.html
- Clark, A. 2006: 'Cognitive Complexity and the Sensorimotor Frontier'. *Proceedings Of The Aristotelian Society*, supp. vol LXXX, pp.43-65.
- Clark, A. 2007: 'Curing Cognitive Hiccups: A Defense of the Extended Mind'. *Journal of Philosophy*, CIV: 4, pp. 163-192.
- Clark, A. 2008: *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press: NY.
- Clark, A and Chalmers, D. 1998: 'The Extended Mind'. *Analysis*, 58: 1, pp. 7-19. Reprinted in P. Grim, (eds) 1998: *The Philosopher's Annual*, vol XXI. Also reprinted in D. Chalmers (ed) 2002: *Philosophy Of Mind: Classical and Contemporary Readings*. Oxford University Press.
- Cosmelli, D and Thompson, E (forthcoming) Embodiment or Envatment? Reflections on the Bodily Basis of Consciousness. To appear in John Stewart, Olivier Gapenne, and Ezequiel di Paolo (eds) *Enaction: Towards a New Paradigm for Cognitive Science*. MIT Press.

- Dennett, D. 1969: *Content and Consciousness*. London: Routledge & Kegan Paul.
- Dennett, D. 1981: 'Where am I?' in D. Dennett *Brainstorms*. Harvester Press: Sussex.
- Dennett, D. 1991: *Consciousness Explained*. Little Brown: Boston.
- Dennett, D. 1996: *Kinds of Minds*. Basic Books: NY.
- Dretske, F. 1981: *Knowledge and the Flow of Information*. Blackwell: Oxford, UK.
- Eliasmith, C. 2008: 'Dynamics, Control, and Cognition', in M. Aydede and P. Robbins (eds) *Cambridge Handbook of Situated Cognition*. Cambridge University Press: NY, pp. 134-154.
- Fodor, J. 1998: *In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind*. Cambridge, MA: MIT Press.
- Hardcastle, V. 1998: 'The puzzle of attention, the importance of metaphors'. *Philosophical Psychology*, 11: 3, pp. 331-352.
- Haugeland, J. 1998: 'Mind Embodied and Embedded', in J. Haugeland *Having Thought: Essays in the Metaphysics of Mind*. Harvard University Press: Cambridge, MA, pp. 207-240). Originally appeared in *Acta Philosophica Fennica*, 58, 1995, pp. 233-267 (a special issue on Mind and Cognition edited by L. Haaparanta and S. Heinamaa).
- Hurley, S. 1998: *Consciousness in Action*. Cambridge, MA: Harvard.
- Hurley, S. 2008: 'Vehicles, Contents, Conceptual Structure, and Externalism'. *Analysis*, 58: 1, pp. 1-6.
- Hurley, S. forthcoming: 'Varieties of externalism', in Richard Menary (ed) *The Extended Mind*. Ashgate.
- Hurley, S. and Noe, A. 2003: 'Neural Plasticity and Consciousness'. *Biology and Philosophy*, 18, pp. 131-168.
- Hutchins, E. 1995: *Cognition In The Wild*. MIT Press: Cambridge, MA.

- Hutchins, E. In press: 'Material Anchors for Conceptual Blends'. *Journal of Pragmatics*.
- Kelso, J. S. 1995: *Dynamic Patterns*. Cambridge, MA: MIT Press.
- Kelso, J.S. 2002: 'The Complementary Nature of Coordination Dynamics: Self-organization and Agency'. *Nonlinear Phenomena in Complex Systems*, 5: 4, pp. 364 – 371.
- Lamme, V. 2006: 'Towards a true neural stance on consciousness'. *Trends in Cognitive Sciences*, 10: 11, pp. 494-501.
- Lamme, V. and Spekreijse, H. 1998: 'Neural synchrony does not represent texture segregation'. *Nature*, 396, pp. 362-366.
- Lungarella, M. and Sporns, O. 2005: 'Information self-structuring: key principles for learning and development'. *Proceedings 2005 IEEE Intern. Conf. Development and Learning*, pp. 25-30.
- Marcas, V., Zihl, J., and Cowey, A. 1997: 'Comparing the visual deficits of a motion blind patient with the visual deficits of monkeys with area MT removed'. *Neuropsychologia*, 35: 11, pp. 1459-1465.
- Martin, M.G.F. 2002: 'The Transparency of Experience'. *Mind and Language*, 17, pp. 376-425.
- Martin, M.G.F. 2007: 'The Limits of Self-Awareness'. *Philosophical Studies*, 120, pp. 37-89.
- McGinn, C. 1991: *The Mysterious Flame: Conscious Minds in a Material World*. New York: Basic Books, rpt 1999.
- Menary, R. 2007: *Cognitive Integration: Attacking The Bounds of Cognition*. Palgrave Macmillan.
- Millar, A. 2007: 'What the Disjunctivist is Right About'. *Philosophy and Phenomenological Research*, LXXIV: 1, pp. 176-198.
- Millikan, R. 1984: *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.

- Noë, A. 2002: 'Is the Visual World a Grand Illusion?', in A. Noë (ed) *Is the Visual World a Grand Illusion?*. Imprint Academic: Thorverton, UK, pp. 1-12.
- Noë, A. 2004: *Action in Perception*. Cambridge, MA: The MIT Press.
- Noë, A. 2006: 'Experience without the head', in T.S. Gendler & J. Hawthorne (eds) *Perceptual Experience*. Oxford University Press: NY, pp. 411-434.
- Noë, A. 2008: 'Magic Realism and the Limits of Intelligibility: What Makes Us Conscious?', in J. Hawthorne (ed) *The Philosophy of Mind*. Blackwell: Oxford, pp. 457-476.
- Noë, A. and Thompson, E. 2004a: 'Are there neural correlates of consciousness?'. *Journal of Consciousness Studies*, 11, pp. 3-28.
- Noë, A. and Thompson, E. 2004b: 'Sorting out the neural basis of consciousness'. Authors reply to commentators, *Journal of Consciousness Studies*, 11, pp. 87-98.
- O'Regan, J.K. 1992: 'Solving the "real" mysteries of visual perception: the world as an outside memory'. *Canadian Journal Of Psychology*, 46: 3, pp. 461-488.
- O'Regan, J. K. and Noë, A. 2001: 'A sensorimotor approach to vision and visual consciousness'. *Behavioral and Brain Sciences*, 24: 5, pp. 883-975.
- Prinz, J. 2009: 'Is Consciousness Embodied?', in P. Robbins and M. Aydede (eds) *Cambridge Handbook of Situated Cognition*. Cambridge: Cambridge University Press, pp. 419-436.
- Putnam, H. 1975: 'The Meaning of "Meaning"', in K. Gunderson (ed) *Language, Mind and Knowledge*. Minneapolis, MN: University of Minnesota Press. Reprinted in H. Putnam *Mind, Language, and Reality: Philosophical Papers vol.2*, New York: Cambridge University Press.
- Rowlands, M. 2003: *Externalism: Putting Mind and World Back Together Again*. Acumen/McGill-Queen's University Press.
- Rowlands, M. 2006: *Body Language: Representing in Action*. MIT Press.

- Serruya, M., Hatsopoulos, N.G., Paninski, L., Fellows, M.R., & Donoghue, J. P. 2002: 'Immediate, Real-Time Use of a neurally-based control signal for movement'. *Nature*, 416, pp. 141-142.
- Singer, W. 2003: 'Synchronization, binding and expectancy', in M. Arbib (ed) *The Handbook of Brain Theory and Neural Networks* (Second edition). Cambridge, MA: The MIT Press, pp. 1136-1143.
- Smith, M. L., Gosselin, F. & Schyns, P. G. 2007: 'From a face to its category via a few information processing states in the brain'. *NeuroImage*, 37, pp. 974-984.
- Smith, P. 1984: 'Could We Be Brains in a Vat?', *Canadian Journal of Philosophy*, 14.
- Stafford, T. and Webb, M. 2005: *Mind Hacks*. O'Reilly, CA.
- Sutton, J. 2002: 'Porous Memory and the Cognitive Life of Things', in Darren Tofts, Annemarie Jonson, and Alessio Cavallaro (eds) *Prefiguring Cyberculture: an intellectual history*. MIT Press and Power Publications, pp. 130-141.
- Thiele, A. and Stoner, G. 2003: 'Neural synchrony does not correlate with motion coherence in cortical area MT'. *Nature*, 421, pp. 366-370.
- Thompson, E., and Varela, F. 2001: 'Radical embodiment: Neural dynamics and consciousness'. *Trends in Cognitive Sciences*, 5: 10, pp. 418-425.
- Tribble, E. 2005: 'Distributing Cognition in the Globe'. *Shakespeare Quarterly*, 56: 2, pp. 135-155.
- Van Gelder, T., & Port, R. 1995: 'It's About Time: An Overview of the Dynamical Approach to Cognition', in R. Port & T. v. Gelder (eds) *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, MA: MIT Press, pp. 1-44.
- Varela, F., Thompson, E., and Rosch, E. 1991: *The Embodied Mind*. MIT Press: Cambridge, MA.
- Wheeler, M. 2005: *Reconstructing the Cognitive World*. MIT Press: Cambridge, MA.

- Wheeler, M. In Press: 'Minds, Things and Materiality', in C. Renfrew and L. Malafouris (eds) *The Cognitive Life of Things*. Cambridge: McDonald Institute for Archaeological Research.
- Wilson, R. A. 2004: *Boundaries of the Mind: The Individual in the Fragile Sciences—Cognition*. Cambridge University Press: Cambridge, UK.
- Wilson, R. and Clark, A. 2008: 'How to Situate Cognition: Letting Nature Take Its Course', in M. Aydede and P. Robbins (eds) *Cambridge Handbook of Situated Cognition*. Cambridge University Press: New York, pp.55-77.