



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Enjoying the Spread

Citation for published version:

Ward, D 2012, 'Enjoying the Spread: Conscious Externalism Reconsidered', *Mind*, vol. 121, no. 483, pp. 731-751. <https://doi.org/10.1093/mind/fzs095>

Digital Object Identifier (DOI):

[10.1093/mind/fzs095](https://doi.org/10.1093/mind/fzs095)

Link:

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

Document Version:

Peer reviewed version

Published In:

Mind

Publisher Rights Statement:

© Ward, D. (2012). Enjoying the Spread: Conscious Externalism Reconsidered. *Mind*, 121(483), 731-751. [10.1093/mind/fzs095](https://doi.org/10.1093/mind/fzs095)

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.



Enjoying the Spread: Conscious Externalism Reconsidered

Abstract: A variety of recent ‘enactivist’ proposals argue that the material basis of conscious experience might extend beyond the boundaries of the brain and nervous system and into the environment. Clark (2009) surveys several such arguments and finds them wanting. Here I respond on behalf of the enactivist. Clarifying the commitments of enactivism at the personal and subpersonal levels and considering how those levels relate lets us see where Clark’s analysis of enactivism goes wrong. Clark understands the enactivists as attempting to provide hypotheses about the subpersonal mechanisms underlying experience according to which those mechanisms contingently include portions of the environment. But understanding enactivism instead as involving a relational conception of experience at the personal-level, with apparent implications for the location of the subpersonal mechanisms of experience, allows us to make better sense of the enactivist arguments, and make the case for conscious externalism.

Where are the boundaries of the physical states underlying conscious experience? Clark (2009) assesses three attempts to argue that those boundaries do not coincide with those of our neural and nervous systems, and concludes that we have no good reason to accept hypothesis of the Extended Conscious Mind (henceforth ECM): the hypothesis that the physical basis of the conscious mind extends into the environment. The arguments Clark considers are all drawn from ‘enactive’ views of experience, according to which capacities for perceptual experience are essentially dependent on capacities for embodied agency in an environment. Clark’s conclusion is that accepting the good points of these enactivist arguments does not commit us to accepting ECM. In what follows, I argue that this assessment rests on a misinterpretation of the fundamental insight of enactivism. Enactivism is a thesis about experience, a personal-level phenomenon. Enactivists argue for their view of experience both by appeal to subpersonal-level considerations concerning the dynamic nature of information processing and by appeal to personal-level considerations of epistemology and phenomenology. However, these different sources of support are not always clearly distinguished by enactivists, nor is the interplay between the enactivist’s claims at the personal and subpersonal levels always made clear. This lack of clarity is what allows Clark to return a negative verdict on the commitment of enactivism to ECM. I do not aim here to show that enactivism or ECM is correct. A precondition of assessing enactivism is gaining a clear view of its theoretical commitments, and that is the goal of this paper. My suggestion will be that Clark is right to note that enactivism, construed as providing hypotheses about the subpersonal bases of experience, does not commit us to accepting ECM. And yet, the uniform acceptance of ECM among enactivists (Noë (2004), Noë and Thompson (2004), Thompson (2007), Hurley (1998, 2008)) is not due simply to confusion on their part. Once we understand the enactivists as urging a specific

personal-level conception of experience we can see why ECM follows from their views.

1.

What is required for the truth of ECM? Clark identifies the hypothesis of the extended mind (Clark and Chalmers (1998), Clark (2008)) with the claim that 'the local material vehicles of some aspects of human cognition may, at times, be spread across brain, body and world' (Clark 2009 p.966). A possible way this could occur is if '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' (Ibid.). But whilst advocates of the extended mind (such as Clark) standardly focus on the claim that such a possibility is realized by the material bases of standard dispositional beliefs, advocates of ECM make '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' (p.967). We determine whether this is true of the vehicle of some particular conscious experience by determining whether the subpersonal basis of the experience is exhausted by events and properties within the organismic boundaries of the experiencing subject. So I understand ECM as the claim that *the subpersonal basis of some conscious experiences can include events and properties outwith the organismic boundaries of the experiencing subject*. I take it that there is no substantive difference between this and Clark's gloss of ECM as the analogous claim about the *material basis* of experience. I focus on the claim about the subpersonal rather than the material basis only to avoid commitment to the idea that the best way of characterizing that subpersonal basis will be in physical terms (as opposed to, say, the language of ecological optics).

With this in mind, it is important to note that there is nothing about the notion of the subpersonal that restricts it to the neural in particular. As Hurley notes:

The causal relations between nervous systems and environments are intricate and continuous. There is nothing specially oomphy about causal relations inside the skin, or inside the head, nothing specially capable of pushing or shoving. So there is nothing causally mysterious or inhospitable to materialism or naturalism or realism about relational states of persons. And there is no magical causal boundary around persons. Viewed subpersonally, they are in principle transparent to causality. (Hurley 1998, p.336)

Suppose I get a beer from the fridge. An appropriate personal-level description of what is involved in my achieving this might be 'reaching into the fridge to get a cold beer'. When shifting to a subpersonal description of the same event, there is no automatic reason to ignore the bits of my body and the world involved in this reaching and restrict our attention to the neural. An appropriate subpersonal

description of the event might involve fine-grained details of the trajectory of my arm, hand, and later the beer, through time and space that do not figure in the personal-level description of my intentional activity.¹ Similarly, if the appropriate personal-level characterization of one of my perceptual episodes is 'looking into the fridge at a cold beer', then it is at least an open question whether a subpersonal characterization of that episode should restrict itself to a specification of my neural states, or also include the portions of the environment involved in my looking and seeing. This is fully compatible with agreeing that the neural portions of the subpersonal characterization of my relationship with the environment are certain to be of special interest in the projects of understanding action and perception. But giving some kind of explanatory priority to the neurally-mediated portions of our relationship with the environment doesn't rule out the possibility of non-neural bits of the world figuring in an appropriate subpersonal characterization.

How do we determine the type of personal and subpersonal characterization of a perceptual episode that is 'appropriate'? In particular, how might characterizations at one level constrain those at the other? Initially, we might suppose that what we say at one level is independent of what we say at another – after all, we know that it is fallacious to draw conclusions about features at one level from features of our talk about the other without supporting arguments (Millikan 1984, Dennett 1991). Nonetheless, it is natural to suppose that there is an interplay between what we say at one level and at the other (see and cf. McDowell 1994a, Bermudez 1995, Davies 2000). Our characterizations of (for example) what a perceiver sees at the personal level are constrained by our views concerning the features of the environment to which her information-processing resources might intelligibly support sensory access; we will not attribute infrared vision to a perceiver if there is no intelligible way of their transducing information about infrared frequencies in their environment. And we could not make sense of attributing mental states to an organism if we found them to be completely hollow inside, or filled with a homogenous jelly (McDowell 1994a). Likewise, what we plausibly know about the personal-level can prompt us to revise our conception of how things are subpersonally. Were we to find a perceiver who unfailingly manifested perceptual sensitivity to infrared light in their thought, behaviour and introspective reports despite the apparent impossibility of their transducing information about infrared frequencies, then this should prompt us to revise our descriptions of them at the subpersonal level, perhaps arriving at a new view of where their sensory transducers are, or what sort of information can pass through them.

¹ Of course, there could be reasons why we might wish to restrict our subpersonal focus to the neural – say, if we were interested only in the mechanisms of fine motor control behind the action. As we will see, descriptions at one level of explanation can be influenced and constrained both by our conception at the other level, and by our explanatory interests.

With these preliminaries in place, I can outline the view of the relationship between enactivism and ECM for which I will argue here. Enactivism offers us a particular personal-level conception of experience – one according to which experiences are essentially episodes of interaction between a subject and parts of the world. We noted above that there is no a priori reason to restrict the subpersonal to the neural, and that a specific conception at the personal-level can make it natural to adopt a complementary conception at the subpersonal-level (and vice versa). Since the enactivist understands experiences as interactive relationships between subject and environment, it is natural for them to resist a subpersonal conception restricted to the neural activity of the organism engaged in such interaction. Indeed, as we will see, the picture of the subpersonal that would result from doing so threatens to make it unintelligible how their conception of things at the personal level could be correct. If we can give an exhaustive subpersonal characterization of my experiential state solely in terms of how things are with me internally then how can that state, when viewed at the personal level, be essentially a dynamic relationship between me and my environment? Hence enactivists think that ECM follows fairly trivially from their position. This also gives us an inkling of why enactivists think their conception of experience and their commitment to ECM is important. We can contrast the enactive view of experiences as essentially world-involving interactions with a view of experience as in principle independent of the environment of the experiencer. On such a view, perceptual experience is not intrinsically an encounter with the environment, but a ‘mere affectation of a person’s subjectivity’ (Putnam 1999, p.11) that, if all has gone well, is appropriately caused by environmental factors. Upon this conception, experience is

confined to a tract of reality whose layout would be exactly as it is however things stood outside it, and the notion of a vantage point on the external world is now fundamentally problematic. (McDowell 1986, p.241)

Clark notes that rejecting the enactivist arguments for the reasons he identifies ‘leaves open the possibility of other (one might say ‘more metaphysical’) arguments that might be thought to support something akin to ECM’ (Clark 2009, p.968), such as the naïve realist view that perceptual experiences can have external objects as constituents. My contention here is that the enactivists are most plausibly read as providing arguments of this kind, in favour of a conception of experience as a skill-mediated relationship between subject and world. The enactivist’s motivation for embracing ECM, it will emerge, is that rejecting ECM appears to commit us to an incompatible ‘internalist’ conception of experience, of the kind Putnam and McDowell describe. Once we are internalists at the subpersonal level, it becomes unintelligible how the enactivist’s personal-level conception of experiences as modes of interaction with one’s environment could be the case.

2.

This view of the relationship between enactivism and ECM emerges most clearly when we consider the enactivist argument that Clark finds most puzzling and least convincing – Noë’s considerations about virtual representation.² One way of approaching these considerations is by noting that, at any given instant, it seems that we experience far more than is allowed by various facts about the information-processing that underpins our perceptual capacities.³ For example, both the resolution and colour-sensitivity of our parafoveal vision is very poor. Yet we experience the world as uniformly coloured and sharply defined throughout, not as monochromatic and blurry in all portions outside our fovea. Noë draws the moral that the relevant parafoveal detail and colour really is present to us in our experience of the world, but that this presence is, in a sense, *virtual*. It obtains in virtue of our current sensory contact with the world *plus* our implicit understanding of the movements we might make to bring other portions of the world into view, and our implicit expectations about the sensory results of such movements. An upshot of this view seems to be that subject’s experience has its rich content only *potentially*, corresponding to the subject’s potential for skillful interaction with worldly objects and structures. However, Clark rightly notes, in assessing ECM we are not interested in the material vehicles of various *potential* experiential contents, but in those which account for the content that experiences actually have. So whilst it might be plausible that experience has its rich content only potentially, this looks irrelevant to establishing ECM. At this point, Clark reads Noë as responding that the content of experience is virtual all the way in – that we cannot factor experiential content into occurrent and potential parts:

Pick any candidate for an occurrent factor. Now consider it. It is structured too; it has hidden facets or aspects too; it is present only in potential. (Noë 2006, p.421)

But, Clark objects, this reading is obscure; if read simply as the claim that we can always attend to the objects of perception in more detail it is clearly true, but looks irrelevant to establishing ECM. If read as the claim that content and character of experience isn’t determined by ‘a snapshot moment of neural activity, but a process extended in time’ (Clark 2009, p.974), then opponents of ECM can simply appeal to temporally evolving neural processes rather than static ones when identifying the basis of experience. He concludes that the argument ‘is either flawed or simply too opaque to carry the weight required.’ (Ibid. p.975)

Understanding enactivism as primarily concerned with claims about the nature of personal-level experience rather than about mechanical underpinnings provides a diagnosis both of why these considerations should be thought to

² Here, and with the enactivist arguments that follow, I restrict myself to a very brief sketch. For a more detailed synopsis, see Clark 2009, section 4. For Noë’s remarks on virtual presence and its relationship to enactivism, see Noë 2004, 2007, 2008, 2009.

³ See Clark 2002 for a useful review.

support ECM, and why Clark finds it difficult to see their relevance. Clark reads these considerations as concerning the way in which temporally extended interaction with the environment might figure in an account of the subpersonal underpinnings of experience, and is puzzled as to how these considerations are intended to prevent us from viewing the relevant environmental interactions as mere causal inputs to and outputs from the neural processes that realize the experience. But Noë's aim in his remarks on virtual presence is not to argue directly for a particular picture of the subpersonal underpinnings of experience, but to motivate a particular account of experiences at the personal level as essentially interactive encounters with the world. We should not interpret his remarks on presence as entailing that the detail of the world, or the occluded parts of objects are present to us in experience 'only potentially', or in some impoverished sense. Whole objects and detailed scenes are embraced in experience, despite facts about the limits of sensory uptake that might appear to make this problematic. Noë's enactive theory is intended as an account of how this is so:

Perceptual experience is always an encounter with objects, but it is also always an encounter with objects that is shaped by one's situation of perspective. The challenge of the theory of perception is to appreciate how perception can be, in this way, an encounter with how things are, when the nature of things necessarily exceeds what can be taken in at a glance. (Noë 2008a, p.691)

According to enactivism, experience is a mode of temporally extended skillful interaction with the world. Though the perceptual information present in our sensory uptake from a particular time and place is impoverished in various ways, the enactivist holds that our understanding and expectations concerning the ways in which changes in this information march in step with changes in our relationship to the environment allow us to transcend these limitations and be related to the world, not just its impacts on our sensory receptors, in experience. We do visually experience the whole cat, not just its unoccluded portions or its facing side, because our practical understanding couples with our current sensory contact with our environment so as to relate us to the worldly three-dimensional cat in an appropriate way. Similarly, the colour and detail of the portions of the world currently in my parafoveal vision can figure in my experience, since my understanding and expectation concerning my sensory contact with the world is such as to put me in an appropriate relation to those parts of the world. For the enactivist, this type of relationship with the world, suffused with practical understanding, is what perceptual experience *essentially is*. Once we read Noë's remarks on virtual presence as aimed at supporting this conclusion, we can make sense of why he accords them such significance in his work (Noë 2004, 2007, 2008a, 2009). We can also see why Clark's suggested readings of Noë on this point miss the mark. It is not being suggested that an evolving sequence of internal experiences, caused by the external environment, is responsible for our perceptual contact with the world. Instead, the enactivist's claim is that an

evolving perceptual relationship with the world, implicating sensorimotor understanding, *is* our experience of the world.

The above also suggests how the enactivist can sidestep an apparent problem for the argument from virtual representing to ECM on which Clark does not dwell. Noë's claims about virtual representation are intended to alleviate an apparent tension between the richness of experiential content and the paucity of the information carried by our sensory relation to the world at a time. However, as Clark notes (pp.967-8), ECM is a claim about the *vehicles* of experience, not about its contents. And there is no straightforward inference from the fact that the content of an experience concerns an extra-neural state of affairs to the conclusion that the vehicle of that experience is also, in part, extra-neural. But the above gloss on Noë's remarks shows that he needs no such inference. Rather, he offers a personal-level conception of experience as an interactive relationship with the environment as a plausible explanation of how the content of experience can outstrip that to which we have sensory access at a time. Such a conception appears to have a straightforward bearing on ECM (we will take up the question of whether this appearance is misleading in section 4); if some experiences essentially involve an interactive relationship between the subject and the distal objects and properties with which they are in actual and potential sensory contact, then the subpersonal bases of those experiences must encompass both parties in that relationship – hence they must outrun the organismic boundaries of the subject. Hence, Noë's remarks on virtual representation are intended to support ECM by providing a personal-level conception of experience that appears to entail it, not by providing subpersonal evidence in its favour.

We can extend this diagnosis to Clark's dissatisfaction with the first set of enactivist arguments he considers, the arguments from variable neural correlates. Enactivists appeal to various bodies of empirical work, such as use of tactile-visual sensory substitution systems (TVSS), adaptation to perceptual distortions, and other instances of neural plasticity, that suggest that the neural correlates of a given experience can vary both across different individuals and across time in a single subject.⁴ Clark rightly notes that all these results look compatible with environmental interactions playing a merely instrumental role, providing the causal input for tuning and training the neural resources that ultimately support experience. Again, it is true that such results do not force ECM on us when we view them as mere observations about dynamic causal interactions with the environment. But enactivists primarily appeal to such results in order to demonstrate the truth of enactivism, not of ECM. Reflection on such cases is supposed to make plausible the idea that experience, at the personal level, essentially is a mode of interaction with the environment. TVSS users only enjoy vision-like experiences through their TVSS devices when they attain a level of proficiency at the mode of environmental interaction (and an attunement to its distinctive sensorimotor contingencies) that is characteristic of vision. The

⁴ See Clark 2009, section 3 for a fuller summary. See Hurley 1998, ch.9, Hurley and Noë 2003, and Noë 2004, 2009 for the original arguments.

experiences of wearers of goggles which invert or distort the information that the environment delivers to the eye return to normal to the extent that they learn to put that information to use in the mode of interaction characteristic of vision. Events internal to the perceiver can differ markedly, but such cases are intended to provide evidence that sameness of interactive relationship with the environment goes hand in hand with sameness of experience. This is supposed to support a conclusion not about the subpersonal, but about what experience, for the enactivist, essentially is – a skill-mediated relation to the world.

I have suggested that in each of the above cases it is a mistake to read the enactivist as concerned with the direct provision of hypotheses about the subpersonal basis of experience. Clark is right to note that their arguments look implausible when read this way. Instead, I suggest, we should read them as suggesting a conception of conscious experience not as something caused in us by the world, but as a specific way of being related to the world. If we endorse such a conception, it is natural to think that a subpersonal description of that relationship should include entities external to skin and skull. When we view experience as an interactive relationship with parts of the world, there is no temptation to suppose that the relevant parts of the world should simply drop out when we shift our focus to the subpersonal. Note that we can certainly ask what makes the sort of relationship with the environment to which the enactivist appeals possible, and it is beyond doubt that processes and structures in the head will play a privileged role in this explanatory project. But the indispensability of an appeal to internal processes and structures in an enabling explanation of how an organism is capable of being related to its environment in a certain way is not something the enactivist calls into question. And granting that internal factors play this crucial enabling role is fully compatible with the enactivist's characterization of what experience *is* – a skill-mediated coupling between perceiver and environment. The enactivist's point is that when we look at the material basis, or subpersonal description, of this coupling (rather than the facts about an organism's structure that enable it to form a part of such a coupling), we find a dynamic flow of energy and information that spans organism and environment, rather than being shorn off at the organism's sensory surfaces.

3.

Our diagnosis of the ease with which Clark rebuts the enactivist's arguments has been that he reads the enactivist as defending ECM via the direct provision of subpersonal hypotheses that vindicate it, rather than via the provision of reasons to accept a picture of experience at the personal level that appears to entail it. Given this, it is natural that the aspect of enactivism Clark finds most persuasive with respect to establishing ECM, and the one to which he devotes most attention, is the idea that the subpersonal underpinnings of experience are a looping dynamical mess, spanning brain, body, and world.⁵ His initial response

⁵ See Clark 2009 section 5 for a summary of the considerations about dynamic entanglement. Enactivist views on the topic can be found in Thompson and

to this suggestion is to wonder whether the neural cross-sections of such dynamic loops might not be carved off, and support experience on their own, were the relevant temporal sequence of neural events somehow to take place within them. For example, might not a disembodied brain in a vat enjoy perceptual experience if somehow ushered through the right sequence of neural events? But, he then suggests, the enactivist might claim that certain evolutions of neural states might *only be possible* in the presence of certain environmental scaffolding. If this were so, then the embedding vat that provides the inputs and registers the outputs of the disembodied brain would need to be functionally equivalent to a normal embedding environment, the necessity of whose contribution the defender of ECM wishes to establish. Clark arrives at this interim conclusion:

At this point in the dialectic, a deep stalemate beckons. Just because OUTER drives BRAIN, and BRAIN depends (let's assume) on OUTER to step through the signature sequence of states that support some specific experience, that doesn't 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. (Clark 2009, p.983)

But things look better on our rendering of the dialectic. The crux of the enactivist's case for externalism is not (contra Clark) the supposition that some sequences of neural events can only unfold as part of a tight environmental coupling. It is the provision of a conception of experience that removes the temptation to focus on the neural portions of a perceiver's dynamic informational entanglement with the environment when setting out the subpersonal basis of a perceptual state. If we accept the enactivist's views about experience, then there is no threat of stalemate. For, as we have seen, the enactivist is not merely offering a picture of an extended dynamical subpersonal basis of experience and suggesting that it is as good as any other. Rather, the subpersonal picture that secures ECM is made natural by their particular view of experience at the personal level, for which they have given independent arguments (such as those concerning virtual representation and perceptual plasticity). Contra Clark, the hypothesis that experience is underpinned by dynamic feedback loops that include portions of the environment is not the enactivist's best argument for ECM – it is rather *a statement* of ECM, motivated by their personal level conception of experience.

However, Clark ends his paper with an interesting alternative proposal for breaking the stalemate he sees:

Varela 2001, Noë and Thompson 2004, Noë 2008, and Cosmelli and Thompson (forthcoming). Susan Hurley's work (especially Hurley 1998, 2008) contains the best and most extensive discussion of the reasons for and consequences of viewing persons as dynamic singularities at the subpersonal level.

[I]t 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 requires millisecond precision (for evidence for this conjecture, see the review in Singer (2003))? ... In such cases the external environment may well *matter* insofar as it drives the neural systems, but the key effects that enable and explain the quality of the felt experience may then be occurring at time-scales that are only possible within the neural apparatus itself. (Ibid. p.984-5)

The proposal is that ultra high-bandwidth information transmission might only be possible in a neural medium, and that insofar as we have reason to believe that such transmission is essential for consciousness, we have principled reason to privilege the neural in our account of consciousness. It is useful for us to compare how the relevance of this observation about intra-neural information-processing and its possible role in binding looks from each of the perspectives on experience we have been considering. If we construe experience as a 'mere affectation of our subjectivity', in principle independent of the state of the environment, it seems plausible that these considerations tell us something about the subpersonal basis of experience. The idea that facts about our subjectivity are essentially independent of facts about the environment, though the environment can causally influence our experience, finds support in the fact that certain sorts of informational transactions that might be thought crucial for subjectivity can only take place internally. And the influence of our conception of the subpersonal upon our conception of the personal can run both ways. If our subpersonal conception of experience is of a piecing together of disparate bits of information via a kind of processing necessarily restricted to our brains, it becomes difficult to see how experience at the personal level could be, in its nature, a way of being in touch with the external world, rather than pure subjectivity that is, if all goes well, appropriately causally related to the environment.

What if we instead adopt the enactivist conception of experience, according to which experience is essentially a coupled interaction between perceiver and world, involving the exercise of sensorimotor knowledge? A subpersonal construal of this skilled interaction might involve ambient energy flowing from object to perceiver, the temporally extended modulation of this flow as the perceiver's head, eyes and body move relative to their environment, and perhaps the operation of neural resources that mediate the perceiver's grasp of the interrelations between environmental stimuli, their sensory uptake, and their actual and possible movements through that environment. Whilst neural resources would figure crucially in such a subpersonal specification, the conception of experience as essentially a mode of interaction between perceiver and world leaves no possibility of restricting our corresponding conception at the subpersonal level to a neural cross-section of the dynamic flow of information between active perceiver and perceived scene. With this in mind, we can

appreciate why the enactivist will see Clark's point about bandwidth and binding as orthogonal to their case for ECM. The considerations about bandwidth demonstrate, at best, that certain intra-neural events are a necessary part of the subpersonal basis of experience. But the enactivist can and should agree with this – we noted above that enactivists should not deny the essential role of neural mechanisms in enabling the kind of relationship between perceiver and environment that is, on their view, constitutive of experience. What is required to rule out ECM is a reason to believe that neural states alone are *always sufficient* for the subpersonal basis of experience. And from the perspective of the enactivist, Clark's point about bandwidth looks simply irrelevant to establishing this claim. I have suggested that ECM follows naturally from the enactivist's conception of what experience is. So, from the enactivist's perspective, casting doubt on ECM requires casting doubt on their conception of experience. And as we have just seen, the point about the necessity of neurally-mediated high-bandwidth interactions does not do this – it can be perfectly naturally accommodated by the enactivist.

4.

I have been suggesting that a commitment to a conception of experience as skilled interaction with the environment is the hallmark of enactivism, and that the enactivist's endorsement of ECM stems from this commitment, not directly from the arguments Clark considers. But isn't this tangential to the real issue? It is not immediately apparent that Clark's discussion of experience rules in or out any particular conception of experience and its dependence or otherwise on environmental interaction. Clark is surely interested in the boundaries of the physical machinery, or minimally sufficient supervenience base, of experience *however* construed. So could we not grant that experiences are fundamentally ways of engaging with one's environment whilst holding that the physical conditions that suffice for experience are exhausted by states of the brain and nervous system? This takes us back to the point with which we closed section 1, and upon which other sections have touched – a key motivation of the enactivist's endorsement of ECM is scepticism about whether internalism at the subpersonal level can be combined with externalism at the personal level in this way. Recall that the enactivist has two related reasons for resisting this combination. Firstly, if we conceive experience as a dynamic relationship between subject and environment, the temptation to go internalist when describing experience at the subpersonal level simply will not arise. If our focus at the personal level is on the subject, the environment, and their relationship then there is no reason why that focus should shrink to include only neural properties of the subject when we go subpersonal. Secondly, it is at best problematic and at worst impossible to see how such a subpersonal internalism could be combined with the sort of personal-level externalism the enactivist wishes to endorse. How, upon the subpersonal internalist picture, could experience be anything but a mere affectation of

subjectivity by a world that is only contingently and causally related to it?⁶ Once we view the subpersonal basis of experience as walled off from the environment by skin, skull and sensory transducers, as exhausted by physical states that are essentially independent of everything external to the subject, then how could experience at the personal level be something that essentially encompasses external objects and properties?

An analogy can help to bring out the point. Suppose we hold a relational view of knowledge, according to which knowledge is essentially a relationship between a subject and the worldly state of affairs known. Episodes of knowledge thus depend not just on properties of the subject that can be specified independently of the state of the world and their standing in it (call these ‘internal properties’ of the subject), but also on the way the world is. Given such a conception of knowledge, if we are interested in delineating the material events and processes that underpin an episode of knowing – in giving a subpersonal characterisation of a personal-level state of knowing – then we must look further than the internal properties of the knower. Whilst those properties can guarantee that the subject is just such as to be related appropriately to the world, they cannot guarantee that the appropriate relation does in fact obtain. For that, the requisite bits of the world must play their part. On the conception of knowledge we are considering, knowledge depends on that relation obtaining – on the world playing its proper part. So in restricting our subpersonal descriptions to internal properties of the subject, and thus ignoring the compliance (or otherwise) of the world that determines whether the relation obtains, we leave it underdetermined whether we have described part of the subpersonal basis of an episode of knowing. Our conception of knowledge at the personal level thus informs and constrains what

⁶ Millar (2007) argues that it is possible to make sense of veridical experiences as encompassing environmental structures and properties whilst retaining a conception of experience as essentially independent of states of the environment. For Millar, this is secured by holding that veridical experience involves the exercise of perceptual recognitional capacities, where exercises of those capacities are constituted by the successful picking out of their object. If such a view is tenable then we could endorse the internalist conception of experience whilst agreeing with the enactivist that veridical experience essentially involves relatedness to one’s environment (the relatedness, on this view, being secured by the successful exercise of a recognitional concept). I cannot hope to provide an adequate discussion of Millar’s view, or these issues, in this short piece. But it seems to me that the enactivist cannot accept the picture Millar recommends. Enactivists have principled reasons for holding that dynamic informational entanglement with the environment only issues in experience when enactive or sensorimotor understanding is exercised (see Noë and O’Regan 2001, Noë 2004, 2009). If an enactivist account of perception appeals to recognitional capacities of the kind Millar favours, then the account of such capacities is likely to be given in terms of such enactive or sensorimotor understanding (see Noë (2009)). So, for the enactivist, experience and the exercise of those capacities cannot come apart in the way Millar suggests.

we can intelligibly say at the subpersonal, and vice-versa. Once we adopt a relational conception of knowledge, a complete characterization of the subpersonal underpinnings of an episode of knowledge must make reference to more than the internal properties of the knower. Conversely, if we suppose that, at the subpersonal level, episodes of knowledge are underpinned solely by internal properties of the knowing subject then we rule out a relational conception of knowledge at the personal level.

The same considerations apply to the enactivist's conception of experience. For the enactivist, episodes of experience depend not just on properties of the perceiver that can be specified independently of the world and their standing in it, but on the dynamic relationship between the perceiver and their environment. So whilst the internal properties of the subject might underpin the subject's being just such as to be related to their environment in the appropriate way, those properties alone cannot determine whether the relation obtains. If we adopt the enactivist's conception of experience as a relationship between perceiver and environment then the subpersonal underpinnings of experience must include more than the internal properties of the subject, since those properties leave it underdetermined whether the requisite relationship obtains. Conversely, holding that the subpersonal basis of experience consists solely in internal properties rules out the enactivist's relational conception of experience at the personal-level.

We are now in a position to see just what the enactivist should say about the case of the disembodied brain that we touched on in the previous section. For all we have said, is it not still intuitive that if one brain is in an identical series of physical states to another over some suitable length of time then it will support the same experience, regardless of how or whether they are embedded within an agent and an environment? As we saw in the last section, Clark agrees with proponents of ECM that a standard appeal to a brain in a vat will not do here, since it seems that a properly envatted brain must be dynamically entangled in its environment in a way that is functionally isomorphic to its normally embodied and embedded twin. But if a disembodied twin of my brain went through an identical sequence of physical events by some improbable quantum coincidence (see Block's (2005) review of Noë 2004), is it not reasonable to suppose it would have an identical sequence of experiences to me?⁷ The right response for the enactivist is simply: No. A lone man in an empty room performing a sequence of movements identical to those of a man dancing a tango is not tangoing, for that requires a certain kind of relationship with a partner and with music. An overturned car spinning its wheels is not driving, for driving requires a certain kind of relationship between a car and the road. For the enactivist, experience

⁷ As Clark notes (p.981), granting this possibility – that a disembodied twin of my brain could support experience – would not amount to a refutation of ECM. Even if the material vehicles of conscious experience were restricted to brain-bound events in such quantum improbability cases, it might still be that those vehicles extend beyond the brain in the case of real subjects as they interact with their environments.

essentially consists in an interactive relationship with one's environment. A disembodied brain going through a sequence of physical states in limbo is not experiencing, for it is in no meaningful relationship with its environment. In the first two cases, the man and the car (respectively) are certainly doing their bit. We have described the cases so that had they been embedded in their environments in different ways, nothing more would have been required of them in order to qualify as tangoing or driving. But (given appropriate conceptions of what it is to tango, or to drive) we can grant this without there being any embarrassment or mystery in our denying that the lone man was tangoing, or that the upturned car was driving. Similarly, given the enactivist's conception of experience, we can grant that the physical states of the disembodied brain exactly match my own – that it is just such as to, when appropriately embodied and embedded, enable the sort of relationship with the environment that is constitutive of experience – whilst denying that it supports conscious experience.

We find, I think, just this position in the work of McDowell, who also believes that experience essentially consists in a skillful relationship with the world. He stresses that:

Without any threat at all to the enormous power of cognitive science to enable us to explain our mindedness (in one sense of 'explain our mindedness', not constitutive explanation), we ought to be able to see that the sheer fact that a brain is going through the motions that an embodied brain goes through when a person thinks or experiences is by itself no ground for supposing that there is a mind in there. (McDowell 1994a, p.201)

And, in dealing with the suggestion that it must at least *seem* to a disembodied brain that it is having experience:

Here it is to the point to respond: you might as well suppose it seems to an electronic calculator that things are thus-and-so. (Ibid. p.201)

Our discussion of enactivism gives us the resources to properly understand this remark. For McDowell, understanding arithmetic requires standing in the space of reasons – being in a position with respect to others and to the facts such that you can give appropriate reasons for, and respond appropriately to criticism of, the things you say and do. The calculator stands in no such relations (though it can help people to stand in such relations), and so possesses no understanding. Similarly, for the enactivist, conscious experience requires a specific relationship between the subject and the environment – a dynamic relationship between the perceiver and perceived, mediated by sensorimotor understanding. According to the enactivist, brains (whether embodied or disembodied) stand in no such relations (though they crucially enable people to stand in such relations to their environment) and so do not, by themselves, support conscious experience. This is why, once we adopt the enactivist view of experience, the basis of experience extends beyond the brain. If we endorse enactivism, then ECM must follow.

5.

As noted at the outset, my aim has been to gain a clearer view of enactivism's nature and theoretical commitments rather than to argue directly for the position. In the brief discussions of virtual representation and neural plasticity we glimpsed some of the motivations for such a view. And in connecting our position on ECM to McDowell's views, via our discussion of the disembodied brain, we suggested another motivation. McDowell opposes the Cartesian conception of our minds as seats of pure subjectivity, blocked off from the environment by our sensory transducers, urging instead that we understand mindedness in terms of the relationships between the subject, the environment and others. According to the understanding of enactivism I have recommended, the enactivist urges just the same. McDowell's historical and epistemological reasons for breaking from the Cartesian conception are also reasons in favour of the view of experience the enactivist adopts. The relationship between enactivism and direct realism is a potentially rich and fascinating topic, which I have only hinted at here. But to even see a relationship worth investigating, we must first have a proper understanding of enactivism. If we are to make sense of the enactivist's position and their reasons for endorsing ECM, we must understand enactivism as first and foremost a thesis about what experience at the personal level is – an interactive relationship with one's environment, mediated by sensorimotor knowledge. While enactivists often appeal to facts at the subpersonal level to support their view (such as facts about neural plasticity and dynamic entanglement), their thesis is pitched at the personal level. Critical discussions of enactivism are often hampered by a failure to appreciate this fact. Without the enactivist's conception of experience in mind, enactivist uses of empirical evidence can seem opaque, unmotivated or question-begging. As well as clarifying the relationship between enactivism and ECM, I hope this discussion has laid some groundwork for a better understanding of enactivism and its potential contribution to our understanding of mind and its place in the world.

References

Bermudez, J.L. 2000: 'Personal and Subpersonal: A Difference Without a Distinction'. *Philosophical Explorations*, 3, pp. 63-82.

Block, N. 2005: 'Review of Alva Noë, *Action in Perception*'. *Journal of Philosophy*, 5, pp. 259-272.

Clark, A. 2002: 'Is Seeing all it Seems? Action, Reason and the Grand Illusion'. *Journal of Consciousness Studies*, Vol. 9, No. 5/6, pp. 181-202.

— 2008: *Supersizing the Mind: Embodiment, Action and Cognitive Extension*. New York: Oxford University Press.

— 2009: 'Spreading the Joy? Why the Machinery of Consciousness is (Probably) Still in the Head'. *Mind*, 118, pp. 963-993.

Clark, A. and D. Chalmers 1998: 'The Extended Mind'. *Analysis*, 58, pp. 7-19.

Cosmelli, D. and E. Thompson 2010: 'Embodiment or Envatment? Reflections on the Bodily Basis of Consciousness'. In Stewart, Gapenne, and di Paolo, eds. 2010, pp. 361-386.

Davies, M. 2000: 'Interaction Without Reduction: The Relationship Between Personal and Subpersonal Levels of Description'. *Mind and Society*, 2, Vol. 1, pp. 87-105.

Dennett, D 1991: *Consciousness Explained*. Boston: Little Brown.

Gendler, T.S. and J. Hawthorne (eds) 2006: *Perceptual Experience*. New York: Oxford University Press.

Hawthorne, J. (ed) 2008: *The Philosophy of Mind*. Oxford: Blackwell.

Hurley, S. 1998: *Consciousness in Action*. Harvard: MIT Press.

—2008: 'Varieties of Externalism'. In Menary 2008, pp. 101-154.

Hurley, S. and A. Noë 2003: 'Neural Plasticity and Consciousness'. *Biology and Philosophy*, 18(1), pp. 131-168.

McDowell, J. 1986: 'Singular Thought and the Extent of Inner Space' in his 1998b, pp. 228-260. Originally published in Pettit and McDowell 1986.

—1994a: 'The Content of Perceptual Experience'. *Philosophical Quarterly*, Vol. 44, No. 175, pp. 190-205.

—1994b: *Mind and World*. Cambridge, MA: Harvard University Press.

—1998a: 'Having the World in View: Kant, Sellars and Intentionality'. *Journal of Philosophy*, XCV, pp. 431-491.

—1998b: *Meaning, Knowledge and Reality*. Cambridge, MA: Harvard University Press.

—2007: 'What Myth?' *Inquiry*, 50:4, pp. 338-351.

Menary, R. (ed) 2010: *The Extended Mind*, Cambridge, MA: MIT Press.

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. 2004: *Action in Perception*. Cambridge MA, MIT Press.

—2005: 'Real Presence'. *Philosophical Topics*, 33(1), pp. 235-264.

—2006: 'Experience Without the Head'. in Gendler and Hawthorne 2006, pp. 411-34.

—2008a: 'Reply to Campbell, Martin, and Kelly'. *Philosophy and Phenomenological Research*, LXXVI, 3, pp. 666–673.

—2008b 'Magic Realism and the Limits of Intelligibility: What Makes Us Conscious?' In Hawthorne 2008, pp. 457-76.

—2009: 'Conscious reference'. *The Philosophical Quarterly*, 59 (236), pp. 470-82.

Noë, A. and J.K. O'Regan 2001: 'A Sensorimotor Theory of Vision and Visual Consciousness'. *Behavioural and Brain Sciences*, 24, pp. 939-1031.

Noë, A. and E. Thompson 2004: 'Are There Neural Correlates of Consciousness?' *Journal of Consciousness Studies*, 11, pp. 3-28.

Putnam, H. 1999: *The Threefold Cord: Mind, Body and World*. New York: Columbia University Press.

Schellenberg, S. 2008: 'The Situation-Dependency of Perception'. *Journal of Philosophy*, 2, pp. 55-84.

Stewart, J., O. Gapenne, and E. di Paolo (eds) 2010: *Enaction: Towards a New Paradigm for Cognitive Science*. Cambridge, MA: MIT Press.

Thompson, E. 2007: *Mind in Life*. Cambridge, MA: Harvard University Press.

Thompson, E., and F. Varela 2001: 'Radical embodiment: Neural dynamics and consciousness'. *Trends in Cognitive Sciences*, 5(10), pp. 418-425.