

Draft of a paper published in the *Journal of Consciousness Studies*, special issue on embodied, embedded and extended accounts of phenomenal consciousness, edited by M. Silberstein and A. Chemero, 22: 3-4, 129-47, 2015.

## **Not What it's Like but Where it's Like: Phenomenal Consciousness, Sensory Substitution and the Extended Mind**

**Michael Wheeler**

### **1. Locating the Topic**

According to the hypothesis of *extended cognition* (henceforth ExC), the machinery of mind sometimes extends beyond the skull and skin. To put things another way, the defining claim of ExC is that, contra the neuro-centrism of most modern cognitive science and most modern naturalistic philosophy of mind, the parts of the physical world that instantiate or implement cognitive states and processes are sometimes spread out over the brain, the non-neural body and the beyond-the-skin environment. More precisely still, ExC is the view that there are actual (in this world) cases of intelligent thought and action, in which the material vehicles that realize the thinking and thoughts concerned are spatially distributed over brain, body and world, in such a way that certain external (beyond-the-skull-and-skin) factors are rightly accorded cognitive status. In this final formulation of the view, the term 'cognitive status' is really just a place-holder for 'whatever status it is that we standardly grant the brain, in cognitive science and naturalistic philosophy of mind, when talking about the causal wellsprings of intelligent thought and action', although ultimately more would need to be said about exactly what that means.<sup>1</sup> Eye-catching examples of external elements that advocates of ExC often take to have cognitive status (in the relevant sense) include smartphones, tablets and at least some instances of wearable computing, but, in the end, nothing much hangs on such feats of contemporary technological wizardry. Less fancy items such as notebooks (the old-fashioned kind), tally sticks and abacuses would, under the right circumstances, do just as well. The phrase 'under the right circumstances' is, of

---

<sup>1</sup> The first-stop presentation of ExC is by Clark and Chalmers (1998). Clark's own more recent treatment may be found in (Clark 2008). For a field-defining collection that places the original Clark and Chalmers paper alongside a range of developments, criticisms and defences of the notion of extended cognition, see (Menary 2010). The third statement of ExC that figures in the opening sentences of the present paper is my own preferred formulation, one that I have been using, with minor variations, for a while (see e.g. Wheeler 2010, 2013, 2014). As far as I can tell, my way of expressing the view doesn't have any alarmingly idiosyncratic features, and so should be broadly acceptable to all the main protagonists in the debate.

course, a writhing nest of philosophical serpents, and most of the debate over ExC might be cast as a dispute over whether those serpents can be pacified. Nevertheless, in what follows, I am going to pretend that that debate is not only over, but that the right side has won. In other words, I am going to work on the assumption that ExC is true: cognitive extension sometimes happens. What do we fight about now?

The answer, I submit, is phenomenal consciousness, that much-discussed (in philosophy anyway) *what-it's-like-ness* of experience – what it's like for me to see red, taste a vindaloo, or, as I'm doing right now, listen to Fripp and Eno. One's first reaction to this suggestion might be that the proposed flashpoint of phenomenal consciousness is, in reality, nothing more than a damp squib, at least for advocates of ExC. After all, given a broad enough interpretation of the term 'cognition' – one that, in effect, makes it interchangeable with the term 'mind', which is how these terms have typically been used in cognitive science – phenomenal consciousness is just one aspect of cognition, so surely if we have a mandate for extended cognition, we have a mandate for extended phenomenal consciousness. But any such thought would be much too quick. Recall that the truth of ExC requires only that *some* of the material realizers of cognition extend beyond the skin. This opens up a potential gap in relation to phenomenal consciousness, since phenomenally conscious states and processes might be (or might be among) the class of psychological phenomena whose material realizers aren't ever extended.

So much for the 'in-principle' situation. What do the authorities say? Well, as it happens, even some of the most enthusiastic champions of ExC start to sweat and shuffle their feet at the mention of extended phenomenal consciousness. And some go beyond mere nervousness. Andy Clark (one of the original architects of ExC, see note 1) has criticized a number of arguments in the ExC literature that, in one way or another, are aimed at establishing the existence of extended phenomenal consciousness (e.g. arguments such as those found in Hurley and Noë 2003, Noë 2004, and Thompson and Varela 2001). In addition, Clark has presented his own argument to the effect that, assuming the truth of (what he judges to be) our best current science of how consciousness happens, phenomenal consciousness remains a defiantly inner, neurally realized phenomenon (Clark 2009). Clark's declared position, then, is that, as science stands, ExC is true, but only of certain unconscious or non-conscious cognitive states and processes, states such as dispositional belief. The details of Clark's various arguments will not be the focus of attention today.<sup>2</sup> I

---

<sup>2</sup> For what it's worth, my own view (un-argued-for here) is that Ward's (2012) criticisms of Clark's paper are broadly correct. Roughly, Ward argues that Clark misinterprets his opposition, because he (Clark) takes them to be arguing directly for the existence of subpersonal extended material vehicles of consciousness, when in truth they are arguing for a constitutive, personal-level account of the nature of conscious perception, although one that (Ward suggests) has implications for the

am recording his conclusion only to register the fact that it is possible to be a supporter of ExC while remaining an internalist about the material vehicles of phenomenal consciousness. Of course, as has already been intimated, not everyone in the ExC camp agrees with Clark's internalism about consciousness – otherwise he wouldn't have had anyone to criticize. However, even those more radical souls who maintain that the physical realizers of phenomenal consciousness are, at least sometimes, extended beyond the skin nevertheless tend to accept that the matter is not automatically settled by the same arguments and considerations as those that (we are assuming here) settle the case for other extended cognitive phenomena. For example, Alva Noë, in giving voice to what he takes to be the received thinking about such things, notes that '[i]t always seemed that there were obstacles to thinking that consciousness (in contrast with cognition) could extend beyond the limits of the skull' (Noë 2004, p.219). One of Noë's own proposals for how to get over such apparent obstacles is discussed below. Right now the point that matters is this: the word on the street is (a) that phenomenal consciousness presents the advocates of ExC with an extra hurdle to be cleared, and (b) that further work is required to ascertain whether or not they can clear it.

The new debate, then, is over the where-it's-like-ness of what-it's-like-ness. Let's call the externalist claim on the table, the *hypothesis of extended phenomenal consciousness* (henceforth ExPC). According to ExPC, there are actual (in this world) cases of phenomenal consciousness in which the material vehicles that realize the phenomenally conscious states and processes concerned are spatially distributed over brain, body and world, in such a way that certain external (beyond-the-skull-and-skin) factors are rightly accorded whatever status it is that we standardly grant the brain, in cognitive science and naturalistic philosophy of mind, when talking about the causal wellsprings of phenomenal consciousness. In other (simpler) words,

---

vehicular level. In terminology due to Susan Hurley that I'll introduce briefly later in this paper (see also footnote 5 below), Clark may be said to confuse what-quality (constitutive) and quality-enabling (vehicular) levels of explanation. According to Ward, Clark's misinterpretation of the structure of the pro-extended-consciousness case in question standardly results in his criticisms falling short of their intended targets. This criticism seems right to me, although Ward and I will disagree about what precisely is entailed about the material vehicles of consciousness by the personal-level account on offer (Wheeler, in preparation). In a further volley against Clark, Ward proceeds to argue, also correctly in my assessment, that Clark's positive case for internalism regarding consciousness (a case that turns on the claim that a certain capacity that science tells us is currently found only in brains is required for consciousness) establishes, at best, only that the material vehicles of consciousness must always include certain neural processes. It doesn't establish what the vehicle internalist conclusion needs, which is that such processes are always sufficient for consciousness.

if ExPC is true, although the material vehicles that realize conscious phenomenal experience include, and maybe necessarily include, neural elements, they are not restricted to such elements. There will be cases in which those material vehicles also include not only non-neural bodily elements, but also elements located beyond the skull and skin.

To keep our thinking straight here, it is worth emphasizing the fact that ExPC, as we are conceiving it, does not entail that the worldly object or state of affairs that one is phenomenally conscious *of* is part of the relevant material vehicle. The objects of phenomenal consciousness may remain external to the vehicle of such consciousness, even if the vehicle is now partly external with respect to the brain and body. Here is an illustrative example. Let's assume that the sequence of movements that will deliver me to some previously unvisited destination is being assembled interactively, via the use of a navigation app on my mobile phone. If certain conditions are met, then the physical circuitry in my mobile phone will be part of the material realizer of my navigation-related cognitive states or processes. That's just boring old ExC. But now let's assume that various aspects of my activity are accompanied by states or processes with phenomenal character (some sort of what-it's-like-ness dimension), and that it is at least partly in virtue of such states or processes that I am conscious of various salient environmental features, such as roundabouts, junctions and landmarks. According to ExPC, if certain conditions are met, then the physical circuitry in my mobile phone will sometimes count as part of the material vehicle of my phenomenal-quality-realizing, environment-disclosing states or processes, whereas the roundabouts, junctions and landmarks that I'm conscious of won't. One might state the general lesson like this: what ExPC requires is the (sometimes) extended character of the physical machinery that enables us to be conscious – in a phenomenal-quality-endowed way – of worldly objects or states of affairs.

With ExPC in better view, it is time for us to narrow our focus to just one particular set of considerations in the vicinity. As mentioned earlier, there are a number of arguments for ExPC already lodged in the literature, but there's not enough room in a single paper to be exhaustive and one has to start somewhere. So I intend to concentrate on what I shall call *arguments from sensory substitution*. As the moniker suggests, arguments of this form endeavour to establish ExPC on the basis of a careful analysis of a well-documented psychological phenomenon, namely that of sensory substitution. As it will concern us here, sensory substitution (examples soon) occurs when technological augmentation enables one sensory modality, for instance touch, to support the kind of environmental access and interaction ordinarily supported by a different sensory modality, for instance vision. There are important and difficult questions, both empirical and conceptual, raised by sensory substitution, regarding, for example, the individuation of the senses, and the character of crossmodal perceptual organization, questions that I shall touch on only

in passing (for more detail, see e.g. Auvray and Myin 2009, Macpherson 2011, Farina 2013, Stiles and Shimojo, forthcoming). Here I am interested only in assessing whether cases of sensory substitution provide evidence for the truth of ExPC. To this end, I shall examine two arguments to the effect that they do, one due to Noë (2009) and one due to Kiverstein and Farina (2012). My final conclusion (spoiler alert) will be that both of these arguments fall short. If the phenomenon of sensory substitution provides evidence for the truth of ExPC, it is not because of the considerations marshalled by these authors in the papers in question.

## 2. The Tools for the Job

Before turning to sensory substitution itself, I shall introduce a piece of conceptual apparatus, partly because it helps us to achieve a better understanding of ExPC, but also because I shall appeal to it – in what is, as we shall see, an aberrant fashion – at a key moment in what follows. The apparatus in question comes from Susan Hurley (2010). It's what she calls the *autonomy metaintuition for phenomenal qualities*. I shall call it simply the *autonomy metaintuition*. As she puts it, '[the] autonomy metaintuition for phenomenal qualities is an expression of the intuition that there is an intractable explanatory gap between physical or functional properties and phenomenal qualities' (Hurley 2010, p.104). Hurley introduces the autonomy metaintuition during her discussion of a position that she calls *what-quality externalism*, according to which phenomenal quality (or character) is partly determined by external (beyond-the-skin) factors. For philosophers anyway, perhaps the easiest way to get a grip on what-quality externalism as a position is to see it as the phenomenal-consciousness-related analogue of the more familiar philosophical view that Hurley calls *what-content externalism* (standardly known as just 'content externalism'). According to what-content-externalism, the contents of mental states (e.g. what follows the 'that' clause in a belief attribution such as 'Elsie believes that water is wet') are partly determined by external factors, meaning that two thinkers who are internally identical in every way may nevertheless possess intentional states with different contents merely through being located in different environments.<sup>3</sup> Analogously, the what-quality externalist holds that two thinkers who are internally identical in every way may nevertheless possess conscious states with different phenomenal qualities merely through being located in different environments.

All that said, our primary business here is not with *what* determines phenomenal quality, but with *where* in space the material realizers are that instantiate such

---

<sup>3</sup> Without giving the details, the classic arguments for externalist claims about mental content involve extrapolations from twin-earth-style thought experiments. These thought experiments were designed originally to establish externalism about meanings in the case of language (Putnam 1975).

qualities. Thus our concern is with what Hurley calls *quality-enabling externalism*, which is her term for ExPC, the view that the material realizers of phenomenal qualities sometimes extend beyond the skin. *Quality-enabling externalism* is the phenomenal-consciousness-related analogue of *content-enabling externalism*, according to which the material realizers of mental content extend beyond the skin, which is roughly equivalent to ExC.<sup>4</sup> Inevitably, there are challenging theoretical issues concerning the relations that obtain between what-quality externalism and quality-enabling externalism,<sup>5</sup> but, for our purposes here, only one of those issues matters, and that is the issue of whether the autonomy metaintuition, introduced by Hurley in terms of the what-determines question, applies also to the how-enabled (or where-realized) question. As Hurley (2010, p.148, note 10) argues, the answer is ‘yes’, but to appreciate why, we need to understand exactly what she aims to achieve by introducing the metaintuition in the first place.

If one simply encountered the autonomy metaintuition in isolation, one might be forgiven for thinking that it is no more than a neologism for the infamous hard problem of consciousness, but that would be to under-estimate its strategic worth, because what it achieves is the disentanglement of two intuitions in the vicinity of phenomenal consciousness that are deeply unhappy bedfellows. The key point, then, is that the autonomy metaintuition is widely held, by many scientists and philosophers, alongside a further deeply engrained intuition that, as Hurley (2010, p.115) notes, results in a combination that is ‘puzzling, even paradoxical’. In the what-quality case, the conflicting intuition is that phenomenal quality is determined entirely by internal (inside-the-skin, standardly neural) factors. This leads to what Hurley nicely dubs the *magical membrane problem*.

Why are intuitions favoring what-quality internalism so strong, given the autonomy meta-intuition? If we have no understanding of how phenomenal qualities *could* be explained, why is the conditional intuition so strong that *if* phenomenal qualities can be explained at all, it could only

---

<sup>4</sup> I say ‘roughly equivalent’ because, depending on how one understands the notion of ‘content’, it may well be that not all mental phenomena for which one might seek physical realizers will bear content. This is especially true if content is understood to be representational content, and if one is impressed by, for example, phenomenologically-inspired arguments concerning the non-representational, yet cognitive, character of hitch-free skilled practical activity (see e.g. Wheeler 2005).

<sup>5</sup> One particularly pressing issue will be whether or not one’s view at the what level places constraints in one’s view at the how level. For example, could one be a what-quality externalist and a quality-enabling internalist? In effect, this is the question that I raised earlier, in note 2, in the context of Ward’s response to Clark’s quality-enabling internalism. I don’t propose to engage this issue here, but see (Wheeler, in preparation).

be in terms of internal factors? Why does the internal-external boundary sustain fiercely internalist intuitions about *what, if anything, must* explain phenomenal qualities despite the general admission of bafflement about *how anything could possibly* explain phenomenal qualities, including neural properties? This is what I call the *magical membrane problem*. (Hurley 2010, p.115)

Once the tension between the internalist intuition and the autonomy metaintuition in the what-quality case is exposed, the extension of the latter to the case of quality-enabling externalism (ExPC) is revealed, since the *very same kind of tension* accompanies the realizer-oriented case. As all the fuss about neural correlates of consciousness indicates, cognitive scientists and naturalistic philosophers of mind widely assume that the material vehicles that realize phenomenal qualities are internally located (that's the internalist intuition about vehicles) while expressing bewilderment about precisely how anything physical could realize phenomenal consciousness (that's the autonomy meta-intuition regarding vehicles). And, as we might put it, echoing Hurley's own rhetorical words, if we have no understanding of how phenomenal qualities could be realized in neural elements, why is the conditional intuition so strong that if phenomenal qualities can be realized physically at all, it could only be in terms of internal (neural) factors (that's the vehicular version of the magical membrane problem)? Indeed, one might even argue (although Hurley doesn't) that the enabling case is primary. Perhaps it's *because* we have no understanding of how phenomenal qualities could be *realized* in purely physical states and processes alone that we have no understanding of how phenomenal qualities could be *determined* by purely physical states and processes alone.

Speculations aside, what seems clear is that the autonomy metaintuition applies straightforwardly to the quality-enabling case. And this suggests something rather interesting. One implication of the puzzling nature of the combination of internalist intuition and autonomy metaintuition, in the realm of phenomenal consciousness, is that, in terms of basic metaphysical plausibility, ExPC and orthodox vehicular internalism about phenomenal character are placed on an equal footing. To assume internalism at the outset here would be to walk headlong into the magical membrane problem. But that means that neither quality-enabling internalism nor quality-enabling externalism (neither orthodox internalism about the vehicles of phenomenal consciousness nor ExPC) have the right to be awarded default status in the debate. What this indicates, I think, is that there is an asymmetry between the strategic positions of ExC and ExPC. In the case of cognitive states and processes where no appeal to phenomenal quality is made, it's eminently arguable that some form of internalism about the realizing material vehicles should be treated as the default view, with the burden of proof resting with ExC (Wheeler 2013). But, as I have just explained, no such internalist presumption is warranted in the case of the

vehicles of phenomenal consciousness. That's the force of the autonomy metaintuition.

At root, the reason for this asymmetry is, of course, that although no commentator on cognitive science should be foolhardy enough to conclude that the discipline is complete, or even on a smooth ride to completeness, nevertheless the fact is that no analogue of the autonomy metaintuition afflicts our understanding of the claim that psychological processes such as representation-building, truth-preserving inference, prototype categorization, memory storage, pattern-matching, and the like are realized in purely physical processes alone. From classical computationalism to connectionism to dynamical cognitive science, there are plenty of candidate models out there, and while there may be all sorts of explanatory problems to be overcome, and while some of these models will end up being dismissed as wide of the mark, none of them ushers in a deep sense of bewilderment as to its potential for being realized in the physical world, equivalent to the sense of bewilderment that is often engendered by phenomenal consciousness. So no analogue of the autonomy metaintuition can gain a foothold in that area of cognitive theory. Surprisingly, then, given all that earlier talk of extra hurdles to be cleared, the conclusion of the present line of thought is that ExPC enjoys a strategic advantage over its *prima facie* more plausible cousin, ExC.

That completes my introduction of the conceptual apparatus that we will need, in order to understand and tackle our two target arguments from sensory substitution. So, without further ado, let's turn to those arguments. The first is due to Alva Noë.

### **3. Altered States**

As mentioned earlier, sensory substitution, as it will concern us here, occurs when technological augmentation enables one sensory modality to support the kind of environmental access and interaction ordinarily supported by a different sensory modality. The seminal work in this area is Paul Bach y Rita's (1972; Bach y Rita and Kerel 2002) research on tactile-vision sensory substitution (henceforth TVSS). In this work, blind subjects were equipped with a head- or shoulder-mounted camera that conveyed information, from video images, via the activation of an array of vibrators located on the subject's abdomen or thigh. After a short period of adaptation, those TVSS subjects who actively controlled the information received, either by manipulating their bodies or by manipulating the camera, were able to make reliable judgments about things such as the number, relative size and position of distal objects in three-dimensional space, and to perform actions such as reaching out and picking up objects. TVSS subjects have also been successful at making perceptual judgments involving effects such as looming and object occlusion, and (this time with image-sourced information transmitted via vibrators on the tongue) have

experienced illusory movement effects such as the waterfall illusion<sup>6</sup> (Bach y Rita and Kercel 2002). TVSS is not the only form of sensory substitution. For example, in auditory-vision substitution (e.g. Auvray et al. 2005), video images from a camera are converted into sounds (e.g. the vertical positions of pixels may be correlated with different audio frequencies) and conveyed to the subject via headphones. Again after short periods of adaptation, subjects equipped with such devices have been able to localize, and to recognize the shapes of, distal objects in three-dimensional space.

What is immediately striking about such cases of sensory substitution is that they are examples of a perceptual functionality that is visual, or at least vision-like (see below), in character, even though the proximal sensory interfaces in operation are either tactile or auditory in nature. If, however, sensory substitution is going to provide the foundations of an argument for ExPC, mere functionality won't be enough. Somewhere in the mix there has to be a claim about phenomenal experience. And indeed, as talk of looming and occlusion effects, and of vision-style illusions, already suggests, TVSS subjects are sometimes sources of experiential reports that indicate a transformation in perceptual consciousness. For example, some blind users of sensory substitution systems report visual qualia such as experiences of phosphenes (the seeing of light without light actually entering the eye) (Ortiz et al. 2011). It is because of such reports that Noë (2009 p.57) plausibly understands TVSS to be a 'full-fledged, bona fide example of... [a] transformation in perceptual consciousness'. To be clear, precisely how one should describe the transformation in consciousness involved is, as Noë himself points out, a moot point. Should one say that blind TVSS subjects really *see*, in some non-metaphorical sense of 'see', or do they 'merely' have vision-like (quasi-visual) experiences? This is a problem that propels us in the direction of the thorny question that I mentioned earlier, of how to individuate the senses. I think it's fair to say that the jury remains out on this issue. But, for us, that really doesn't matter, because all Noë needs, in order to construct his argument from sensory substitution to ExPC, is that TVSS engenders a transformation in perceptual consciousness such that, even though the proximal stimuli are tactile in character, *the conscious experience in question is not correctly categorized as one of touch* (Noë 2009, p.62). Given even the limited array of evidence presented above regarding the perceptual sensitivity of TVSS subjects to distal objects, that much seems to be beyond serious doubt. After all, touch is a way of accessing external elements by coming into physical contact with them – surface to surface, as one might say. That is not what's going on in TVSS. And, if any extra argument on this point is needed, it is worth noting that it requires conscious, deliberate effort for TVSS subjects to have tactile experiences of the vibrations taking place on the surface of their skin, rather than the experiences of the distal world that the technological augmentation makes available to them (Noë 2009, p.62).

---

<sup>6</sup> If one stares at a waterfall for a while, and then looks at any stationary rocks at the side of the waterfall, the rocks will appear to be moving upwards.

So how might one argue from the phenomenon of sensory substitution to the truth of ExPC? The pivotal claim in Noë's version of the argument is this: TVSS cases show that '[neural] rewiring isn't necessary for changes in conscious experience' (Noë 2009, p.56). The first thing to appreciate is that the phrase 'changes in conscious experience' refers to those transformations in phenomenal quality in which, as we have just discussed, the transformation in question is from one modality of sensory experience to another. In TVSS, this is a transformation from touch to a modality that is not touch, even if it isn't vision. Of course, if you have poor eyesight, then there is a sense in which putting on the right spectacles will transform your sensory experience, but what the technology does in this instance is provide enhanced inputs that boost performance *within* a modality. The view of the unreconstructed vehicle internalist is not in any way tested by the thought that, when one puts on spectacles, areas of the brain that were contributing to poor vision come to contribute to good vision, without any need for any neural rewiring. But now what about the transformations we care about – transformations from one sensory modality to another? Is neural rewiring necessary in those cases?

For work that at least nudges us towards an affirmative answer to this question, and which additionally gives us a better sense of the *kind* of neural rewiring that attracts Noë's attention, consider experiments on neural plasticity due to Sur and colleagues (Sur et al.1999, described by Noë 2009, pp.53-4). In these experiments, the neurophysiology of new-born ferrets was rewired so that their eyes ended up being connected to the parts of their brains ordinarily used for hearing. The result was that those regions of the ferret brain that standardly process auditory information were recruited for vision. This constitutes a transformation in sensory experience of the right kind (from one modality to another), and it is obvious that an important – one might even be tempted to think, necessary – factor in effecting that transformation was neural rewiring. But, according to Noë, what the case of TVSS tells us is that one shouldn't succumb to temptation here, or at least not if, in doing so, one arrived at a general rule to the effect that rewiring is *always* necessary for the sort of transformation at issue. After all, in TVSS cases, there has been a transformation in phenomenal quality without any ferret-style connecting up of the subject's eyes to the areas of her brain ordinarily associated with touch. Rather, in spite of the visual or vision-like phenomenal consciousness, the touch areas of the TVSS subject's brain continue, as is normally the case, to receive input via tactile stimulation (vibrations on the skin). Why precisely is neural rewiring ruled out as an explanation here? For one thing, TVSS subjects adapt swiftly to their augmenting devices, with the transformation in experiential sensory modality happening far too quickly for the explanation to turn on the kind of substantial neural rewiring that we see in the ferret case. For another, TVSS subjects are typically adult human beings whose brains simply do not have the plasticity exhibited by the brains of new-born ferrets.

So, it looks as if Noë is right: what TVSS cases indicate is that, in the relevant sense, neural rewiring isn't necessary for changes in conscious experience.

This is undoubtedly an intriguing result, but how, according to Noë, does it provide support for ExPC? Spelling things out in a little more detail than Noë himself does, the argument, I think, must go something like this. Begin by assuming an orthodox naturalistic approach to the material realizers of phenomenal consciousness, according to which phenomenal qualities are realized solely in the brain. With this vehicle-internalist assumption in place, and given that no significant neural rewiring has taken place in the TVSS subject's brain, we seem to be left with no naturalistic explanation for the observed transformation in phenomenal character. For under these conditions of explanation, the lack of neural rewiring in TVSS cases means (so the argument goes) that there has been a transformation in the character of conscious experience without there having been any relevant change in the material vehicles which realize that experience. Substance dualism might have something to offer us by way of an explanation here, but naturalistic philosophy of mind is, it seems, silent. A moment's reflection, however, indicates that if we drop the opening assumption, that the realizers of phenomenal quality are wholly internal, then naturalism is back in the game. For if phenomenal consciousness may sometimes be realized partly externally, then it's the addition of some new externally-located realizing factors, in the guise of the TVSS-delivering technology, that promises to eliminate the salient explanatory gap. And that's how the claim that neural rewiring isn't necessary for a transformation in conscious experience supports ExPC. It throws down an explanatory challenge that ExPC can meet, but which quality-enabling internalism can't. Thus, Noë (2009, p.65) concludes, 'the machinery of mind is extended'.

It is worth pausing to note that Noë's own brand of ExPC has some distinctive features. In brief, Noë's view is that perceptual experience is constituted by implicit knowledge of what he calls sensorimotor contingencies – the law-like effects that either my movement or the movement of objects in my sensory field have on the sensory input that I receive (see e.g. Noë 2004, 2009; O'Regan and Noë 2001). Thus, for example, say I see something as a straight horizontal line. According to O'Regan and Noë (2001), my implicit knowledge of how sensory input will change as I move in particular ways is constitutive of that perceptual experience as a case of my seeing a straight horizontal line. Such implicit knowledge includes, for example, my knowledge that as I shift my visual fixation point, the curvature of the line as traced on my retina will change in certain reliable ways, and that, as I look away from the line, its cortical representation will change its shape from (roughly) that of a straight sausage with squashed ends to that of a banana. This approach to perceptual experience plausibly has the resources to explain the fact that TVSS subjects have vision-like experience, since the sensorimotor contingencies in play involve relations characteristic of visual phenomena such as looming and occlusion effects. So far, so

good. It is at this point, however, that an interesting question arises, regarding the relationship between sensorimotor contingency theory and ExPC, namely, is the former inevitably externalist in form, or could it, in principle, be formulated in an internalist register? This is a tricky question (see Silverman 2014 for enlightening analysis and discussion), but I don't propose to pursue it here, because, on my interpretation, Noë's argument from sensory substitution to ExPC may be formulated independently of sensorimotor contingency theory.<sup>7</sup>

Back to the main plot. How persuasive is Noë's argument? To my mind, the weakest link is the transitional conclusion that there has been a transformation in the character of conscious experience without there having been any relevant change in the neural vehicles concerned. Of course, if the only kind of neural change that could count as a relevant change is one that involves ferret-style neural rewiring, then Noë's argument for ExPC would, I think, go through. But, as I am about to suggest, the situation regarding the ongoing neural contribution is not as clear-cut as Noë's treatment suggests. Indeed, there are at least two ways in which other relevant neural changes might have occurred which, in principle at least, might explain the observed transformation in phenomenal consciousness, leaving the internalist free to reject the transitional conclusion. So these alternative scenarios constitute problems for Noë's argument, as I have reconstructed it. The following point helps to bring things into view. In effect, Noë's argument assumes that the realizing neural vehicle of interest in TVSS cases is somatosensory cortex, and somatosensory cortex *alone*. He writes:

Stimulation of the skin gives rise to neural activity in touch areas of the brain (the so-called somatosensory cortex). But for a person who has adapted to the sensory substitution system, activation in somatosensory touch areas gives rise not to the experience of being touched (or at least not only to the feeling of being touched) but to a visual experience of the scene in front of him. (Noë 2009, p.58)

Of course, even considering just the phenomenological evidence available, it's unsurprising that somatosensory cortex remains active during TVSS perception. After all, as mentioned above, subjects are able to shift their conscious attention to the tactile vibrations on their skin and experience them as such. However – and here we are on the doorstep of the first problem for Noë's argument – it is an aspect of the ExPC-justifying set-up, as Noë presents it, that no other region of the brain that might realize the transformed phenomenal quality – a region such as, for example visual cortex – is activated crossmodally *along with* somatosensory cortex. If areas in

---

<sup>7</sup> Further investigation of the issues just raised would, once again, lead us to engage with the kinds of questions introduced (but not explored) in notes 2 and 5 above.

visual cortex are included in the realizer of the post-transformation phenomenal quality, then quality-enabling internalism remains a live explanatory option.

This is obviously a matter ripe for empirical investigation, so what do the existing imaging and stimulation studies tell us? The data is a little messy. Some studies of sensory substitution show activation in visual areas in early blind subjects but not in sighted subjects. However, others show activation in visual areas during recognition and localization tasks in early blind, late blind and (less regularly) blindfolded sighted subjects. Moreover, studies in which parts of the brain are deactivated in order to establish causal links between particular areas and subject performance indicate that sensory substitution devices can recruit visual cortex.<sup>8</sup> Further complications are introduced by the fact that late blind users, who of course have been sighted, may use visualization capacities that activate visual cortex in a top-down way (Stiles and Shimojo forthcoming). There are, of course, many questions of detail to be asked here, especially regarding the organizational nature and the adaptive timescales of the apparent crossmodal plasticity and interaction. The point for us, however, is this: the success of Noë's argument from sensory substitution depends on the empirical bet that visual cortex is not activated alongside somatosensory cortex in sensory substitution subjects. That bet may well be lost.

Still, let's assume that, for some subjects anyway, that bet is won: for those subjects, the realizing neural vehicle of interest is definitely somatosensory cortex alone. Under these circumstances, would Noë's argument from sensory substitution be rehabilitated? I am not convinced that it would. The crucial claim, recall, is that, in sensory substitution subjects, there has been a transformation in the character of conscious experience without there having been any relevant change in the neural vehicles concerned. Noë's benchmark for a relevant change in the neural vehicles is ferret-style rewiring, and it's the thought that this sort of extensive rewiring can't have happened which helps secure the externalist conclusion. But now what about another possible kind of neural change – a fundamental change in the mathematical structure of the neural activation patterns in somatosensory cortex? If a fundamental change in activity patterns in the same neural area, without any actual rewiring, and without any relevant change in the neuro-sensory input channel, can produce a transformation in the modality of the perceptual experience, from, say, touch to vision or quasi-vision, then Noë's argument for ExPC won't go through, because there will be a vehicle-internalist explanation for the target change in experience.

How should one respond to this suggestion? Connectionist research in artificial intelligence and cognitive science has of course made us familiar with the idea that

---

<sup>8</sup> The summary conclusions from imaging and stimulation studies listed here are taken from a recent review paper (Stiles and Shimojo forthcoming), where a large number of references to specific studies may be found.

phase transitions in the activation dynamics of neural-style networks can generate very different systemic behaviour, without any structural reconfiguration of the units and connections (without any neural rewiring). Nevertheless one might remain suspicious of the my tentative claim that such phase transitions may produce a transformation not only in the functional behaviour of the system, but also in the modality of sensory experience. It is at this point, however, that we need to remind ourselves of Hurley's autonomy metaintuition, as applied to the quality-enabling domain. That intuition is that we have no understanding of how phenomenal qualities could be realized in physical or functional elements. Deployed in the service of ExPC, the next stop is the magical membrane problem. But, at this point in the debate, the autonomy metaintuition can equally be deployed in a deviant fashion, that is, in the service of quality-enabling internalism. For if we have no understanding of how phenomenal qualities could be realized in physical or functional elements, why is the intuition so strong that if phenomenal qualities can be realized physically at all, that realization must be specified at the level of neural wiring, rather than at the level of patterns of activation. What this deviant use of the autonomy metaintuition suggests is that we should remain open-minded about the possibility that phase transitions in neural activation dynamics which take place without structural reconfiguration (neural rewiring) might nevertheless produce shifts in the modality of conscious experience. This is, of course, a long way from a knockout blow against Noë's argument from sensory substitution. For one thing, the autonomy metaintuition also urges us to be open-minded about partly external physical realizers of phenomenal consciousness. But it does suggest that the argument fails to deliver a decisive outcome in the dispute between ExPC and quality-enabling internalism.

#### **4. Technology Incorporated**

Before closing, I want to consider, rather more briefly, an alternative argument from sensory substitution, one due to Kiverstein and Farina (2012). This argument turns on a redeployment of Clark's (2008) distinction between the 'use' of a tool and its 'incorporation'. As Kiverstein and Farina (2012, p.36, quoting partly from Clark) explain, the incorporation of a tool, as opposed to its mere use, occurs when 'the brain has been recalibrated so as 'to automatically take account of new bodily and sensory opportunities''. The prototypical example offered of incorporation is Maravita and Iriki's (2004) study of modifications in the body schema (the brain's model of the body in space) during tool use. In this study, macaques are trained to use rakes to retrieve food. As a result of this training, the body schema in the macaque brain is updated so that it counts the rake as 'part of' the monkey's body. (More precisely, bimodal neurons adjust their receptive fields so as to respond to stimuli located at the end of the rake.) According to Kiverstein and Farina, it's precisely this kind of neural recalibration that causally underpins the phenomenological experience in which an expertly manipulated tool is rendered

transparent to agential consciousness. (Recall the famous example of the blind person whose cane is so fluidly integrated into her navigational activity that the phenomenologically fixed boundary between agent and world is shifted from the body-cane interface to the end of the cane.) Kiverstein and Farina observe that this is exactly the kind of experience enjoyed by sensory substitution subjects, whose phenomenological interface with the world is not, under normal circumstances, the body-technology tactile interface, which becomes transparent, but rather the technology-world interface that aligns with their vision-like consciousness. Putting together the pieces of the picture, the authors suggest, quite plausibly, that, during sensory substitution, the subject's neurally-realized body schema will have been modified, so that the augmenting technology is internally represented as being part of the subject's body. The result is that sensory substitution exemplifies the phenomenon of incorporation.

Kiverstein and Farina take incorporation to be sufficient for ExPC. As they put it, where there is incorporation, 'the substrate of the experiences of the user extends to include the device' (Kiverstein and Farina 2012, p.35). So, if sensory substitution is a case of incorporation, and if incorporation is sufficient for ExPC, then sensory substitution is sufficient for ExPC. Let's agree that sensory substitution is a case of incorporation. Is incorporation sufficient for ExPC? The answer, I think, is 'no'. It's simply hard to know why a modification to the internal representational resources of the subject – a modification to the neurally located machinery that determines how the agent categorizes portions of the world as body and not-body – should have the implication that the material vehicles where phenomenal consciousness is realized now include the elements that have been newly represented as body rather than not-body. The inference remains puzzling, I think, even where those neural modifications have the added effect of showing up in experience as the transparency of the technology to phenomenal consciousness. To put things in a mildly imprecise way, what has certainly changed is the experience of where the bodily boundary is located. What hasn't thereby changed is where, in relation to the bodily boundary, experiences are located.

## **5. Concluding Remarks**

In this paper, I have examined two arguments for the claim that the material vehicles that realize phenomenal consciousness extend beyond the skin. Both of these arguments turn on a treatment of the intriguing phenomenon of sensory substitution. My depressingly negative conclusion is that neither argument is successful. Of course, this doesn't show that the hypothesis of extended phenomenal consciousness is false. The decisive supporting argument may, right now, be waiting in the philosophical shadows, poised to leap out on unsuspecting passing internalists. If that's right, however, the argument in question will need more in its kit-bag than sensory substitution, rakes and ferrets.

## Acknowledgments

For useful critical discussion, many thanks to audiences in Stirling and Memphis, and especially to Philip Ebert and Peter Sullivan.

## References

- Auvray, M., Hannequin, S., Lenay, C., and O'Regan, J. K. (2005). There is something out there: Distal attribution in sensory substitution, twenty years later, *Journal of Integrative Neuroscience*, 4, pp.505–521.
- Auvray, M. and Myin, E. (2009) Perception with compensatory devices: From sensory substitution to sensorimotor extension, *Cognitive Science*, 33, pp.1036–1058.
- Bach-y-Rita, P. (1972) *Brain mechanisms in sensory substitution*, New York: Academic Press.
- Bach-y-Rita, P. and Kerckel, S. (2002) Sensory substitution and augmentation: Incorporating humans-in-the-loop, *Intellectica*, 2 (35), pp.287-297.
- Clark, A. (2008) *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*, New York: Oxford University Press.
- Clark, A. (2009) Spreading the joy? Why the machinery of consciousness is (probably) still in the head, *Mind*, 118 (472), pp. 963-993.
- Clark, A. and D. Chalmers, D. (1998) The extended mind, *Analysis*, 58 (1), pp.7-19.
- Farina, M. (2013) Neither touch nor vision: Sensory substitution as artificial synaesthesia? *Biology and Philosophy*, 28 (4), pp. 639-655.
- Hurley, S. (2010) The varieties of externalism, in (Menary 2010).
- Hurley, S. and Noë, A. (2003) Neural plasticity and consciousness, *Biology and Philosophy*, 18, pp. 131-168.
- Kiverstein, J. and Farina, M. (2012) Do sensory substitution devices extend the conscious mind?, in Paglieri, F. (ed.). *Consciousness in Interaction: the Role of the Natural and Social Context in Shaping Consciousness*, Amsterdam: John Benjamins.

- Macpherson, F. (2011) Individuating the senses, in Macpherson, F. (ed.) *The Senses: Classical and Contemporary Readings*, Oxford: Oxford University Press.
- Maravita, A. and Iriki, A. (2004) Tools for the body (schema), *Trends in Cognitive Sciences*, 8 (2), pp.79-86.
- Menary, R. (ed.) (2010) *The Extended Mind*, Cambridge, Mass.: MIT Press.
- Noë A. (2004) *Action in Perception*, Cambridge, Mass.: MIT Press.
- Noë A. (2009) *Out of Our Heads: Why you are not your Brain, and other lessons from the Biology of Consciousness*, New York: Hill and Wang.
- O'Regan J. K. and Noë, A. (2001) A sensorimotor account of vision and visual consciousness, *Behavioral and Brain Sciences*, 24, pp.939-1031.
- Ortiz, T., Poch, J., Santos, J. M., Requena, C., Martínez, A. M., Ortiz-Terán, L., Turrero, A., Barcia, J., Nogales, R., Calvo, A., Martínez, J. M., Córdoba, J. L., and Pascual-Leone, A., (2011) Recruitment of occipital cortex during sensory substitution training linked to subjective experience of seeing in people with blindness. *PLoS One*, 6 (8), DOI: 10.1371/journal.pone.0023264
- Putnam, H. (1975) The meaning of 'meaning', in his *Philosophical Papers, Vol. II : Mind, Language, and Reality*. Cambridge: Cambridge University Press.
- Silverman, D. (2014) *The Sensorimotor Theory of Perceptual Experience*, PhD thesis, University of St. Andrews, <http://hdl.handle.net/10023/5544>
- Stiles N. R. B. and Shimojo, S. (forthcoming) Sensory substitution: A new perceptual experience, in Wagemans, J. (ed.) *The Oxford Handbook of Perceptual Organization*, New York: Oxford University Press.
- Sur, M., Angelucci, A. and Sharma, J. (1999) Rewiring cortex: The role of patterned activity in development and plasticity of neocortical circuits, *Journal of Neurobiology*, 41 (1), pp.33-43.
- Thompson, E., and Varela, F. (2001) Radical embodiment: Neural dynamics and consciousness, *Trends in Cognitive Sciences*, 5: 10, pp. 418-425.
- Ward, D. (2012) Enjoying the spread: Conscious externalism reconsidered, *Mind*, 121 (483), pp.731-751.

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

Wheeler, M. (2010) In defense of extended functionalism, in (Menary 2010).

Wheeler, M. (2013) Is cognition embedded or extended? The case of gestures, in Radman Z. (ed.), *The Hand, an Organ of the Mind: What the Manual tells the Mental*, Cambridge, Mass.: MIT Press.

Wheeler, M. (2014) Revolution, reform, or business as usual? The future prospects for embodied cognition, in Shapiro, L. (ed.) *The Routledge Handbook of Embodied Cognition*, Abingdon and New York: Routledge.

Wheeler, M. (in preparation) Extended phenomenal consciousness: An interim report, for a special issue of the *Southern Journal of Philosophy* collecting together papers from 2014 Spindel Conference at the University of Memphis, on *Alternative Models of the Mind*.