

Phenomenal Intentionality Meets the Extended Mind

Terry Horgan and Uriah Kriegel, University of Arizona

The Monist 91 (2008): 353-380

Introduction

Over the last quarter of the twentieth century, an orthodoxy of sorts had gelled in the philosophy of mind around a kind of *psychological externalism*, the idea that some mental states individuate sensitively to extra-cranial factors. More recently, two trends of thought have departed from this young orthodoxy in opposite directions. On the one hand, a cluster of ideas captured in such phrases as “extended mind,” “embodied cognition,” and “enactive consciousness” proclaim to go further in externalizing the mind; call this the *extended mind outlook*. On the other hand, a constellation of ideas surrounding the notion of “phenomenal intentionality” has ventured to roll back some of the most important aspects of psychological externalism; call this the *phenomenal intentionality outlook*.

In this paper, we examine the relationship between these two opposing trends. Both of us are on record defending the second.¹ Here, we will argue that the phenomenal intentionality outlook can accommodate the letter of the so-called extended mind hypothesis, while utterly renouncing the spirit with which it is often embraced, thus neutralizing the alleged philosophical significance of the extended mind hypothesis. The purpose of this exercise is to show that there is nothing in the letter of the extended mind hypothesis that undermines a more traditional, strongly internalist, broadly Cartesian picture of the mind. The extended mind hypothesis is, to that extent, much ado about relatively little.

In §1, we explicate the extended mind’s letter and spirit, as we conceive of them. In §2, we introduce the main tenets of the phenomenal intentionality outlook, and in §3 we elaborate on one key tenet. In §4, we show how this key tenet drives a wedge between

the letter and spirit of the extended mind. In §5, we consider and respond to a natural objection to our argument from the previous section. We conclude that since the tenets of the phenomenal intentionality outlook jointly constitute a strongly internalist picture of the mind, and since these tenets are not themselves plausibly threatened by the burgeoning body of work conducted in the name of the extended mind, the extended mind hypothesis comports smoothly with a strongly internalist picture of the mind.

1. The Extended Mind Hypothesis and Outlook

Psychological externalism maintains that some mental states individuate partly in terms of factors external to the subject's head. This is because these mental states individuate in terms of their content, and content externalism suggests that some kinds of content individuate externalistically in this way. The extended mind (EM) hypothesis goes further in claiming that the *states* themselves, not only the factors in terms of which the contents of the states individuate, are partly external to the subject.

Psychological externalism has sometimes been carelessly characterized as holding that, say, beliefs ain't in the head. Dretske (1995) renounces this, saying that beliefs are certainly in the head; it is just that what makes them the beliefs they are is not. This is the natural view for an externalist of his ilk. As Yablo (1997) points out, what makes a penny a penny is not in the pocket, but it surely does not follow that pennies ain't in the pocket. Likewise, it does not follow from the fact that what makes beliefs the beliefs they are is sometimes not in the head that these beliefs are themselves not in the head.

Though it does not follow from psychological externalism that beliefs are not in the head, it may yet be independently true that they are not. The EM hypothesis can be thought of as consisting precisely in this extra claim, claiming on independent grounds that, *sometimes*, beliefs (and/or other mental states) are not in the head.

The "independent grounds" in question are a battery of thought-experiments, due mainly to Clark and Chalmers (1998), of which the following is the most intuitively compelling (1998: 14). Inga and Otto each go to the museum, and on their way each forgets where the museum is. Inga consults her memory, and after some mental exertion, realizes it is on Whatever Street. Otto consults his notebook and realizes that it is on

Whatever Street. For some ways of filling in the story about Otto, Clark and Chalmers argue, the two are cognitively indistinguishable. Thus if Otto's notebook meets certain conditions,² one can think of it as part of an unusual realizer of Otto's belief that the museum is on Whatever Street, a realizer distributed between Otto's brain and his notebook. If so, one ought to say that while Inga's belief is entirely in her head, Otto's is distributed between his head and notebook.³

One way to think of the resulting view is as a sort of *vehicle externalism* (Hurley 1998). The familiar version of psychological externalism is founded on externalism about some contents, but it construes the vehicles that carry those contents—i.e., the token states that instantiate the mental state-types with those contents—as intra-cranial.⁴ The EM hypothesis, by contrast, claims that even the vehicles of mentation sometimes individuate externalistically.

It is worth noting that, on at least one construal of the internalism/externalism contrast, the EM hypothesis can also be held in conjunction with content internalism.⁵ This is probably going to be impossible if we construe content externalism as the claim that contents sometimes individuate sensitively to factors outside the head, or as the claim that content properties sometimes do not supervene on the intra-cranial properties of the subject. But if we construe content externalism (ultimately much more plausibly) as the thesis that content properties sometimes do not supervene on vehicular properties, then it becomes possible to combine content internalism with vehicle externalism. The view would be that the content properties of mental states always supervene on those states' vehicular properties, but the vehicular properties of some mental states fail to supervene on the intra-cranial properties of those states' subjects.

The kind of vehicle externalism suggested by the EM hypothesis is thus logically insulated from a number of wider theses. But proponents and sympathizers of the EM hypothesis have often taken it to be pregnant with important implications for one's conception of the nature of mind. Although the letter of the EM view is, as just explained, somewhat logically insular, the view's spirit is much more sweeping.

One way to think of the view's spirit is as undermining the traditional contrast between an inside and outside of mental life. On a philosophically familiar, broadly Cartesian picture of the nature of mind, there is a sort of interiority that constitutes the

individual's mind and that, though certainly not *causally* independent of the rest of reality ("the external world"), is very much *constitutively* independent of it. This broadly Cartesian picture is threatened already by content externalism, inasmuch as the latter holds that, for some mental states, there is *no* constitutive independence from the outer world after all. But the challenge presented to the Cartesian picture by the extended mind view is much deeper and more corrosive. It is that there is no non-arbitrary way to draw any robust, stable inner/outer distinction in the first place. It is not just that mental states individuate partly in terms of what goes on in the external world; it is rather that they are themselves *part* of the external world. The mind and the world are enmeshed and continuous with each other, they flow into one another in an interpenetrating manner, rather than stand opposite each other as two poles of a stable subject/object divide. Thus the Cartesian picture of the mind as somehow strongly internal is wrongheaded not so much because, once one draws the inner/outer distinction, one sees that much of mentality is constitutively dependent upon the outer, but rather because there is no general way to home in on anything distinctly inner for which we might ask whether its content is constitutively dependent on, or independent of, the outer.

Clark and Chalmers are clear that skin and skull are arbitrary places to draw an inside/outside distinction. But the point surely generalizes, and any distinction one would draw is apt to appear equally arbitrary. In the case of the Otto's notebook, the artifact to which Otto's mental state extends is relatively close, in space and time, to Otto's skin and skull. But that is an accidental aspect of the case. Giere (2006) offers a different example, that of the Abell 1689 cluster of galaxies, 2.2 billion light years away, that has been "recruited" by astrophysicists, in a manner akin to Otto's recruitment of his notebook, to function as a gravitational lense through which to see even further into the past. The astrophysicist whose mental state is distributed between her brain and a galaxy mind-bogglingly distant in both space and time underscores the radical nature of the EM hypothesis. But it also brings out the deeper significance of the EM outlook: that there is no principled spatial and/or temporal line between what is mental and what is non-mental, between what is "inside" the mind and what is "outside" it. To be sure, there is a sense in which what is inside the skull is "internal," but its being internal in that sense is

not central to its status as mental, and what is central to mentality has nothing particularly to do with anything “internal” in this sense.

We surmise that many philosophers who wish to resist the EM hypothesis are motivated in large part by a desire to protect this kind of broadly Cartesian picture of the mind. They attack the letter of the extended mind view because they are dismayed by its spirit. Our own strategy here will be different. We will argue that the letter of the view supports the spirit of the view only against some background assumptions that can, and ought to, be rejected. Once these assumptions are supplanted by others, the letter of the view can be accommodated by a strongly internalist, Cartesian conception of the nature of mind.

To avoid confusion, we will reserve the label “extended mind *hypothesis*” for the letter of the view, and use “extended mind *outlook*” for its spirit. In a way, our target in this paper is a conditional, one whose antecedent is the EM hypothesis and whose consequent is the EM outlook. Although we will say something about the antecedent later on (in §4), our focus is on the conditional. For our sense is that once the conditional is denied, the plausibility of the antecedent becomes a less urgent matter.

2. The Phenomenal Intentionality Outlook

Going in the opposite direction from the EM outlook is a set of ideas concerning so-called phenomenal intentionality.⁶ We will refer to the loose cluster of these ideas as the *phenomenal intentionality outlook*. In this section, we expound five central tenets of the phenomenal intentionality outlook, and in the next one a sixth. Although it is only the sixth tenet that bears directly on the central argument of the paper, we canvass the other five here for two reasons. First, doing so provides a more fleshed out contrast between the two opposing outlooks we are considering in this paper. Second, the fact that the central premise of our argument coheres with a cohesive cluster of ideas does provide indirect support for the premise.⁷

The first thesis in the cluster is simply that there exists phenomenal intentionality. The phrase “phenomenal intentionality” denotes a kind of intentionality that phenomenally conscious states exhibit, and moreover exhibit precisely *in virtue* of being

phenomenally conscious states, that is, in virtue of their specific phenomenal character. The first thesis, then, is that there is such an intentionality; that there exist phenomenally conscious states that are intentional and intentional in virtue of their phenomenal character. More precisely, the claim is that there are mental states that instantiate properties whose nature is both phenomenal and intentional: the constitutive phenomenal character of such properties is intrinsically intentional.⁸

The second thesis is that phenomenal intentionality is *narrow*, that is, individuates internalistically. It is quite possible that phenomenally conscious states have contents that individuate externalistically. But those are not the contents they have in virtue of being phenomenally conscious states.⁹ The content they bear in virtue of being phenomenally conscious is purely internalistic—because phenomenal character is itself intrinsic, in the sense that it is not constitutively dependent on what happens “outside the head.” The phenomenally intentional content of your current experience would be shared by any molecule-for-molecule duplicate, as well as by a duplicate brain in a vat with no connections whatsoever to the environment.¹⁰

The third thesis is that phenomenal intentionality is inherently subjective, in that it presents its content to the subject. McGinn (1988) puts it by saying that conscious content is Janus-faced: in addition to the outward-looking face, the fact that it presents the world, it also has an inward-looking face, in that it presents what it does *to the subject*. Georgalis (2006) and Kriegel (2003a) highlight the fact that while non-conscious representation instantiates the two-place relation *x represents y*, conscious representation is essentially a matter of the three-place relation *x represents y to z*. Horgan, Tienson, and Graham (2004), as well as Williford (2005), maintain that phenomenal intentionality is self-presenting, in that its subject is necessarily aware of the precise content of her phenomenally intentional states in a direct and incontrovertible way, and purely in virtue of *being* in that phenomenally intentional state.

The fourth thesis is that phenomenal intentionality is exhibited not only by sensory states, but also by more cognitive, “intellectual” states. This is what Horgan and Tienson (2002) call *cognitive phenomenology*. A direct example of it is offered by Strawson (1994). When two people listen to the news in French, but only one of them speaks any French, there is a difference in what it is like for them to listen to the news.

There is a kind of understanding-experience, whose phenomenology is purely cognitive, that the French-speaker undergoes but the other does not.¹¹

The fifth thesis is that phenomenal intentionality is the only kind of non-derivative intentionality. There is a widely accepted distinction between derivative and non-derivative intentionality. Paradigmatic bearers of the former are traffic signs, flags, and indeed words: they represent what they do only because, and to the extent that, they are interpreted the way they are. Thus they derive their content from the act of interpretation. Presumably, however, some things must have their content in and of themselves, and not because they derive it from some other source. The question is what these things are that boast non-derivative intentionality. Traditionally, it has been common to think that all mental states have non-derivative intentionality. The thesis before us is that only phenomenally conscious states have non-derivative intentionality, while non-conscious states' intentionality is on a par with that of language and other conventional systems of representation.^{12,13}

Arguments for this thesis vary, but most take the form of arguing that phenomenal intentionality has a certain unusual feature, not found in other forms of intentionality, that is needed in order to carry content non-derivatively. Searle (1991, 1992) argues that genuine intentionality requires an aspectual shape, such that representations of the morning star can represent Venus specifically under the morning-star aspectual shape; and that only conscious representations have an aspectual shape in and of themselves and not merely derivatively. Horgan and Tienson (2002) argue that genuine intentionality requires determinate content, such that rabbit representations can determinately represent rabbits, as opposed to undetached rabbit parts, un-disconnected rabbit stages, etc.; and that only conscious representations have determinate content in and of themselves and not merely derivatively.¹⁴ Kriegel (2003a) argues that genuine intentionality requires at least the possibility of a subjective component, such that representations can in principle represent what they do *to someone*; and that only conscious representations have the ability to represent to someone (namely, their subject) in and of themselves and not merely derivatively.¹⁵

These five theses, or the loose cluster thereof, can be taken to constitute the tenets of what we call the phenomenal intentionality outlook. To these we would like to add

here a sixth tenet, one which we have not found in the existing literature, but which coheres very comfortably with the other tenets and to some degree flows quite naturally (though not quite deductively) from the fifth. It is the thesis that phenomenal intentionality in some way captures the core of mentality and is crucial to “the mark of the mental.” We call it the *Phenomenally Intentional Mark thesis*, or PIM for short.

3. Phenomenal Intentionality and the Mark of the Mental

The basic idea behind PIM is that phenomenally intentional states are the only states that are uncontroversially, unquestionably, paradigmatically, prototypically mental. On the view we favor, other mental states count as mental only when, and insofar as, they bear the right relationship to phenomenally intentional states. Furthermore, since this relationship comes in degrees, the mental-state status fades away, with gray areas in which there is no deep fact of the matter as to whether a given state is mental or not. More generally, as far as states that are not phenomenally intentional are concerned, there is never a deep fact of the matter regarding their status as mental states. Rather, states qualify more or less strongly as mental in virtue of some relationship they bear to phenomenally intentional states. Phenomenally intentional states are the only ones that qualify as mental in and of themselves and regardless of any relationship they might bear to any other state, or indeed anything else whatsoever.¹⁶

In this section, we elaborate on and defend PIM. We will consider three questions in particular: (a) what is the relationship in which non-phenomenally-intentional states must stand to phenomenally intentional states in order to qualify as *mental* states?; (b) why think that only phenomenally intentional states constitute the elite of uncontroversially, unquestionably, paradigmatically, prototypically mental states?; and finally, (c) why believe PIM at all?

Starting with (a), there can certainly be different answers to it, but our own view involves an instrumentalist appeal to causal role and cognitive architecture. Consider Dennett’s (1987) intentional stance theory. According to Dennett, organisms are ascribed intentional mental states purely on instrumental grounds: it helps predict and explain their observable behavior. Intentional mental states are thus useful fictions without which it

would be difficult for us to understand, predict, and manipulate the behavior of some creatures, but it is not a mind-independent fact that these creatures are in those intentional mental states.¹⁷ Our own view recommends a somewhat mitigated intentional stance approach to the restricted class of mental states that are not phenomenally intentional.

Thus our view departs from Dennett's in two crucial ways – the mitigation and the restriction – but also retains an important aspect of it. The first departure (the mitigation) is that, for non-phenomenally-intentional states, we do not maintain that regarding them as mental is a useful *fiction*. It is useful, alright, but not necessarily a fiction.¹⁸ The second departure (the restriction) is that our view does not apply Dennett's instrumentalism globally, but only to non-phenomenally-intentional states. Phenomenally intentional states, by contrast, are mental states regardless of whether it is convenient or useful to regard them as such. What we retain in Dennett's view is the notion that, for most mental states, what makes them mental is simply the fact that, if one were to take an intentional stance toward the system in which they occur, it would prove instrumental from that stance to regard the state in question as mental.

Our restricted instrumentalism has two main advantages over Dennett's global instrumentalism. The first advantage is that the global variety faces a fundamental problem that does not affect our restricted variety. Dennett's intentional stance holds that states of a system have intentional content only when, and insofar as, it is instrumental to interpret them to. Thus every state derives its content from some actual or possible act of interpretation. But of course, these acts of interpretation are themselves contentful, and so would have to derive their own content from a further act of interpretation. This appears to lead to infinite regress, however, unless one postulates a kind of "original" intentionality that does not derive from and depends upon interpretation. There needs to be such original, non-derivative intentionality before content can start to be passed around, so to speak. The global intentional system requires anchor points of intentional-stance-free intentional states. On our view, phenomenal intentionality is precisely this sort of original, non-derivative intentionality, and phenomenally intentional states serve as the anchor points for the wider system.

The second problem is that an instrumentalist approach is simply implausible for any states endowed with phenomenal consciousness. A phenomenal experience may be

profitably ascribed to a system, but surely it would qualify as a mental state even if it was not. Imagine a race of creatures all of whose phenomenal states are “proto-epiphenomenal”: their only causal power is to produce an introspective judgment to the effect that they occurred. (We may further suppose that those introspective judgments themselves do not alter the behavior of the creatures.) It would seem that each individual creature has every reason to believe that it has phenomenal states, and that those are mental. This despite the fact that there is no instrumental (explanatory and/or predictive) gain in this supposition. Global instrumentalism, however, entails that the creatures have no reason to regard their phenomenal experiences as mental states, since they do not affect behavior.

There is a question, of course, as to which states are likely to actually qualify as mental under this instrumentalist approach, that is, which states will in fact pass the intentional-stance test (if you will). Our view is that the relevant states are those that are causally integrated in the right way within larger systems that feature phenomenally intentional states. The kind of causal integration we have in mind will underwrite cognitive, broadly inferential connections between the relevant states and phenomenally intentional states. (We are using the term “inference” in its liberal reading here, to cover also sub-personal processes that resemble personal-level inferences.¹⁹) These connections allow for all the system’s states to be jointly construed as collectively interpretable under the intentional stance, while also honoring the non-derivative, stance-independent mental status of phenomenally intentional states.²⁰

Furthermore, the tighter the causal integration, the more strongly the relevant states qualify as mental. Some states have fairly immediate causal/inferential connections to phenomenally intentional states, and those will qualify strongly as mental states. Thus, some early states of the visual system (such as Marr’s 2.5D sketches) are inferentially adjacent, as it were, to conscious visual experiences endowed with phenomenal intentionality. Likewise, some perceptual beliefs are inferred fairly immediately from perceptual experiences with phenomenal intentionality. Such states enjoy such strong connections to phenomenally intentional states that they unproblematically qualify as mental. But other states may have more tenuous connections to phenomenally intentional states, and the more tenuous those connections are, the more tenuous the relevant states’

status as mental. Thus, as we go to earlier and earlier stages of visual processing, the inferential connections to phenomenally intentional visual experiences become more and more distant. Ultimately, some of those states (perhaps LGN states, but certainly retinal states) would not qualify as mental due to this growing “inferential distance.” For some states, the inferential distance would be such that they would fall into a gray area where their mental-state status would be vague. For such states, we want to claim, there is no deep fact of the matter as to whether they are mental states or not. They are what they are: states that are connected to phenomenally intentional experiences in some weak but real way.

Let us move on to (b), the question of why we think that only phenomenally intentional states are prototypical, paradigmatic, unquestionably mental states. We will present two main reasons for this, though there may be others.

The first reason can be explained in two steps. First, it seems to us that phenomenally conscious states are the only ones whose status as mental requires no theory and is justified in an entirely pre-philosophical way. This probably has to do with the fact that only phenomenally conscious states are accessible in a first-person (hence non-theoretical) manner. Second, for familiar reasons, to do with the transparency of experience (Harman 1990) and similar considerations that we will not rehearse here, we hold that all phenomenally conscious states are intentional. It follows that the only states whose mental status is pre-philosophically and non-theoretically uncontroversial and unproblematic are phenomenally intentional states.

The second reason is more subtle, and has to do with the fact that phenomenal consciousness and intentionality are overwhelmingly the most central and distinctive aspects of mentality. On this basis, Rorty (1970) claimed that the Cartesian paradigms of mentality are mental images and occurrent thoughts, which are both intentional and conscious, and that other states qualify as mental because they share with those either their intentionality or their consciousness. Rorty himself took this to imply that there is no underlying unity in mentality, as intentionality and consciousness have nothing to do with each other – it just so happens that paradigmatic mental states exhibit both – and therefore there is no commonality between mental states that resemble the paradigms along the intentional dimension and those that resemble it along the phenomenal

dimension. We reject both parts of Rorty's claim about the implication, of course: we hold that in the paradigmatic cases, intentionality and consciousness have a whole lot to do with each other, inasmuch as the paradigmatic mental states have an intentional content that is fully constituted by their phenomenal character; and we hold that there is a real commonality among all non-paradigmatic mental states, namely, that they are all sufficiently causally/inferentially connected to the paradigmatic states to make them instrumentally treated as mental. But we do accept Rorty's starting-point observation that mental images and occurrent thoughts are the paradigmatic mental states. Since both mental images and occurrent thoughts are phenomenally intentional states, it follows that only phenomenally intentional states are paradigmatic.²¹

Finally, let us address question (c): what is the reason to believe PIM, with its two-layer conception of mental status and its designation of phenomenally intentional states as comprising the elite layer? We have already said why one ought to take phenomenally intentional states to comprise the elite layer, but it remains to be seen why a two-layer conception is warranted in the first place. Again we present two main reasons. The first is that it seems to us highly plausible that this is how the folk conception of mentality actually works. The second is that none of the alternative views about the mark of the mental seem to us at all viable.

Start with the first consideration. The traditional view of concepts is as given by sets of severally necessary and jointly sufficient conditions. But, as Rosch (1973) argued, most folk concepts do not actually behave like that. The folk do not usually categorize objects or events by applying such necessary and sufficient conditions to them. Rather, they categorize them by considering how similar they are to certain designated *prototypes* (or, in a later variation, *exemplars*). For example, the concept of a bird is not such that x is a bird iff x is a feathered biped, but rather, it turns out, such that x is a bird just in case, and to the extent that, x is similar enough to a sparrow.²²

The view we propose is the double proposition that (a) the concept of a mental state both *is* and *ought to be* a prototype concept and (b) the relevant prototypes (or exemplars) both *are* and *ought to be* exclusively phenomenally intentional states, with the purest instances being (most probably) visual experiences and conscious occurrent

thoughts.²³ (In addition, we maintain that what will underwrite the relevant similarity are the causal/inferential connections highlighted above.)²⁴

Note that this double proposition is both descriptive and normative. Its descriptive strand (which, we gather, should be experimentally testable) purports to describe the way the psychologically real concept of mentality actually works. Its normative strand attempts rather to state how the concept of mentality operative in philosophical discussions *ought* to work. The idea is that if it worked this way, it would capture more accurately the real nature of the mental than if it worked otherwise. The upshot, in any case, is that on our view, x is a mental state just in case, and to the extent that, x is similar enough to a phenomenally intentional states (such as visual experiences and conscious thoughts).

The second reason to adopt PIM is simply that it is better than the competition. Discussions of the mark of the mental usually take as their starting point Brentano's (1874) proposal that intentionality is that mark, typically glossing over the fact that Brentano held that only conscious states are endowed with intentionality.²⁵ It is clear, however, that intentionality is *too broad* a criterion for mentality, inasmuch as there are non-mental items (e.g., tokened sentences of natural language) that are intentional derivatively, in a way that ultimately depends on the mental intentionality of language-users past and present.

The suggestion might be modified to claim that it is specifically *non-derivative* intentionality that constitutes the mark of the mental. But in this version the criterion would be *too narrow* if, as we maintain, only phenomenally conscious states have non-derivative intentionality.²⁶ Moreover, it strikes us that phenomenal consciousness ought to figure centrally in any plausible conception of the mind, such that zombies are not truly minded creatures. So any criterion that uses any type of intentionality in a way that is divorced from phenomenal consciousness is inadequate.

In that sense, a move in the right direction is offered by Searle (1992), who maintains that every intentional mental state must be at least *potentially* conscious. The suggestion might be made, on this basis, that the mark of the mental is at-least-potential-consciousness: x is a mental state iff x is at least potentially conscious. However, the suggestion faces two serious challenges. First, it is unclear what potentiality means in this

context, and there is a genuine danger that on every precisification the resulting suggestion will be inadequate. Second, the suggested criterion is too narrow, as it excludes a panoply of apparently mental states that are not potentially conscious, in the sense that they cannot *become* conscious consistently with the laws of psychology. A clear example is provided by states of the dorsal stream of the visual system, which govern on-the-fly visually guided behavior but are never conscious (see Milner and Goodale 1995). But there must be many others: Marr's (1982) 2.5D sketches, blindsight states, states of covert face recognition in prosopagnosia – and this is only in the visual system!

Our sense is that a consciousness-based mark of the mental is necessary, but that the only viable version of this idea is PIM. There is certainly much more that can be said on behalf of PIM, both by way of defense and by way of elucidation. But what we have laid out so far should suffice for the discussion to follow. What we want to argue, in the next section, is that PIM undermines the alleged philosophical significance of the EM hypothesis.

4. The Extended Mind in Light of the Phenomenally Intentional Mark

Observe that some of the above rejected marks of the mental would present a problem for the EM hypothesis. Thus, as Adams and Aizawa (2001) meticulously argue, it is likely that extended states of the sort ascribed to Otto do not have non-derivative intentionality, and therefore do not qualify as mental states by the light of the thesis that non-derivative intentionality is the mark of the mental. So the non-derivative intentionality mark would exclude extended mental states.

Likewise, it is highly implausible to suppose that Otto's (or Otto and his notebook's) extended state is potentially conscious (though of course that would partly depend on what potentiality amounts to). If so, it would not qualify as a mental state by the lights of the Searlean mark.²⁷

We have rejected these mark theses and so do not think that they can serve as an argument against the EM hypothesis. But what they bring out is that the proponent of the EM hypothesis must do more than just argue that there are extended states. She or he

must also show that such states qualify as mental. Doing so would involve commitment to a mark-of-the-mental thesis, or to a disjunction of such theses, and it must also be the case that the thesis, or one of them, is plausible.

In their argumentation, Clark and Chalmers (1998) appeal to the following “parity principle”: if Otto’s extended state plays in Otto’s mental life the same causal-cognitive role that Inga’s unextended state does in hers, then Otto’s extended state should qualify as mental just as much as Inga’s unextended state. The simplest mark thesis that could be extracted from this principle is a broadly functionalist one: *x* is a mental state iff *x* occupies the right causal role.

If this is the mark thesis backing the EM hypothesis, however, it is starkly implausible. A state of the Chinese nation, or a sufficiently gigantic arrangement of beer cans, does not qualify as a mental state, no matter how similar its causal role to that of a genuine mental state (Block 1978).

Nonetheless, we think that, against the background of PIM, a not implausible case might be made for taking Otto’s extended state to qualify as mental. PIM allows mental states of two kinds: (a) phenomenally intentional states and (b) states suitably connected to phenomenally intentional states. Otto’s extended state is clearly not a phenomenally intentional state, as it is cognitively on a par only with a dispositional, non-conscious belief. But it may well be *suitably connected* to phenomenally intentional states, namely, via the kind of causal/inferential connections that underwrite the fitting application of an intentional stance.

Let us stress that we do not wish to endorse the claim that Otto’s extended state *is* in fact suitably connected to phenomenally intentional states; merely that we can envisage a case to that effect being made. Our point is that *if* such a suitable connection can be shown, *then* by PIM’s light we should regard Otto’s extended state as a genuine mental state. Clark and Chalmers (1998) do bring up a number of compelling considerations supporting the instrumentality of treating Otto’s extended state as a mental state. On the other hand, we are also impressed by Rupert’s (2004) battery of considerations suggesting that the utility of doing so is bound to be limited, and that it is accompanied by considerable disutility.

One very reasonable view, and which we are tempted to endorse, is that Otto's extended state is a good example of the kind of gray area mentioned in the previous section, where it is not clear whether or not a state should be regarded as mental, and where ultimately there is no deep fact of the matter as to whether it is. It does seem, after all, that if Otto's extended state passes the intentional stance test, it does so far from spectacularly.

More important for our present purposes is the fact that even if Otto's extended state passed the intentional stance test with flying colors, that would not affect our recommended conception of the nature of mind, as captured by PIM and more generally the phenomenal intentionality outlook. In other words, even if the EM hypothesis could be justified by PIM (and, incidentally, we cannot see how it could be justified in any other way, given how poorly it fares by the light of the aforementioned non-functionalism marks of the mental and how implausible the functionalist mark is), it would not thereby support anything like what we called the EM outlook.

Recall that the main idea behind the EM outlook is that, when it comes to the mind and its place in the world, there is no significant inner/outer distinction to draw, such that some or all of what goes on "inside" the mind might be claimed to be constitutively independent of what goes on "outside" it. In particular, there is nothing about skin and skull that makes them delineate a distinction central to the nature of mind.

We agree that there is nothing special about skin and skull.²⁸ But there is something special about something else, and that is the internal tokening of phenomenally intentional states such as visual experiences.²⁹ So although we share Clark and Chalmers' irreverence toward the skin, rather than go out and away from the skin, we would like to go further *inward*, to the neural correlates of prototypical mental states. For on our view it is in terms of similarity to those states that the scope of mentality is defined.

Suppose that the neural correlate of visual consciousness consists in synchronized neural activity in the visual cortex and the dorsolateral prefrontal cortex (Kriegel 2007a). Then one can take the brain region composed of the visual cortex and the dorsolateral prefrontal cortex (a scattered object) to be the location of the tokening of these strongly prototypical mental states, and therefore to delineate a meaningful inside/outside

distinction that is central to the nature of mind. Since, according to us, the nature of mind is captured most acutely by similarity to these sorts of prototype, there is a very clear interiority here that is preserved regardless of whether Otto's extended state counts as mental.

To bring out the Cartesian flavor of the present picture of mind, consider an envatted duplicate of Otto's brain, or even more minimally, of the part of Otto's brain wherein are tokened Otto's prototypical mental states, i.e., his phenomenally intentional states. The minimal duplicate will share with Otto all these prototypical mental states, and the slightly larger duplicate (the whole brain) most or all of Otto's other mental states. The only states the duplicate will not share are Otto's extended states, which may not be mental states, and if they are, are such only very marginally.³⁰ More perspicuously, the states the duplicate will not share are ones for which there is simply no deep fact of the matter as to whether they are mental. It would seem, then, that Otto and his duplicate will share their unquestionably mental lives, and may or may not share the aspect of their life that is anyway questionably mental. If so, the conditional "if EM hypothesis, then EM outlook" is false, inasmuch as PIM can accommodate the EM hypothesis while rejecting the EM outlook.

It might be objected that, since the whole point of the EM outlook is that there is no substantive inner/outer distinction to be drawn, requiring the proponent of that outlook to provide a mark of the mental by the lights of which extended states qualify as mental amounts to begging the question. According to the EM outlook, the objection goes, there simply may not be a mark of the mental. Indeed, it is most natural for the EM outlook to incorporate a tenet to the effect that there is no mark of the mental.

In assessing this objection, it is important to keep in mind that our immediate target in this paper is not the EM outlook itself, but the conditional "if EM hypothesis, then EM outlook." Our main claim is that the EM hypothesis does not provide strong support for the EM outlook, as it is routinely taken to, since it is fully and smoothly compatible with the diametrically opposed phenomenal intentionality outlook. Given this dialectical situation, the objection under consideration would undermine our argument only if the non-existence of a mark of the mental would fall out of the EM hypothesis.

But that does not seem to be the case. Nothing about the case of Otto and Inga, for example, suggests that there is no mark of the mental.³¹

It might be true that, although the EM outlook is not supported by the EM hypothesis, it *is* supported by the conjunction of the EM hypothesis and the thesis that there is no mark of the mental. But for this reasoning to constitute a compelling case for the EM outlook, some case would have to be made for the thesis that there is no mark of the mental. As just noted, nothing about the EM hypothesis supports the no-mark thesis, so an independent case would have to be mounted. We are unfamiliar with any such case.

Although our primary target is not the EM outlook, but the conditional of which it is the consequent, we do have in mind the plausibility of the EM outlook as a secondary target of sorts. Our thought about the plausibility of the EM outlook is simply this. The EM outlook is antecedently implausible. The only reason one might accept it is because it appears to be supported by the EM hypothesis. So if it turns out, upon examination, not to be supported by the EM hypothesis, then there is really no motivation to adopt it.

5. Phenomenal Vehicle Externalism

The above discussion depends, however, on the substantive claim that phenomenally intentional states are “unextended,” and therefore the intra-cranial nature of their tokenings can be taken to constitute a kind of interiority. Although Clark and Chalmers (1998: 12) appear to concede as much, other proponents of the EM outlook do not. In particular, enactive approaches to perceptual consciousness, of the sort defended by Hurley (1998) and Noë (2004), have sometimes been thought to involve a sort of vehicle externalism about phenomenally conscious states – a *phenomenal vehicle externalism*, if you will. If an astrophysicist’s visual experience can be distributed between her brain and Abell 1689, then there may be no meaningful inside/outside distinction even against the background of PIM. What we will argue in this section, however, is that phenomenal vehicle externalism is implausible. More specifically, the thesis of enactive consciousness is ambiguous: in one reading it supports phenomenal vehicle externalism but is utterly implausible, in another reading it is quite plausible but nowise supports phenomenal vehicle externalism.

The starting point of the enactive approach is the idea that the phenomenology of ordinary conscious experience is mischaracterized when it is construed as a collection of sensations. The actual feel of visual experience is as of a much more dynamic and engaged give-and-take with the environment. It is the feel as of interacting with the world, exploring it, and enjoying a certain know-how of “getting around” in it. Ultimately, the phenomenology of visual experience, for example, is a bundle of potentialities, or dispositions – dispositions to explore and interact with the world.

Consider a visual experience of an apple (Noë 2006). At some brutally sensory level, you are presented in your experience only with the half of the apple facing you, or perhaps just half of the apple’s surface. But surely the overall phenomenology of this experience is as of an encounter with an apple, an entire apple, not half a surface of an apple. In a sense, then, the backside and the inside of the apple are present in the phenomenology. But in *what* sense? In the sense that I have an experience as of being able to walk around the apple and bring its backside into view, or reach and grab it and either turn it to reveal its backside or cut it to reveal its inside. Thus the visceral phenomenology is, in part, as of a cluster of behavioral dispositions toward the apple.

It is often thought that an upshot of this enactive account of perceptual phenomenology is that the supervenience base of phenomenal consciousness is not fully located in the brain (O’Regan and Noë 2001). This would entail that tokenings of phenomenally intentional states are themselves extended. Thus the account seems to construe the very vehicles of visual experience, and more generally phenomenal intentionality, as extended.

This kind of “extended consciousness” view might serve as a basis for an objection to our argument from the previous section. It might be claimed that, even if one embraces PIM, the EM outlook is still correct, because even our designated prototypical mental states are extended. In other words, to defend our rejection of the EM outlook, we must do more than just embrace a view that grounds our conception of the mind in phenomenally intentional prototypes. We must also reject the enactivist claim that phenomenally intentional states are extended.

Our argument against such phenomenal enactivism is straightforward, however. It is that enactivism casts phenomenal consciousness as a *dispositional* property, whereas

consciousness is most certainly not a dispositional property, but an occurrent, manifest one. For someone to undergo a visual experience of a bright green car is not – not only – for her to be disposed in certain ways. There is something very real and categorical going on in her. It may well be that having the experience disposes the subject in certain ways (and that the dispositions are *relational* vis-à-vis the car). But it is not as though the experience *consists* in those dispositions. For, as we stress, phenomenal consciousness is not a dispositional property, but a categorical one.

We think it is obvious that phenomenal consciousness is not merely a dispositional property – obvious in some immediate, first-personal way. But perhaps the following consideration might add further support. Phenomenal consciousness is introspectively accessible. But it would seem that if it were merely dispositional it would not be. Just as fragility and solubility are not perceivable, though they are thinkable, so if phenomenal consciousness were merely dispositional, it would not be introspectible, though it would be higher-order thinkable. We are working here with a conception of introspection as involving something more than just thinking about the introspected—i.e., as involving some sort of quasi-perceptual contact with the introspected. We realize that this is not an entirely uncontroversial conception of introspection, but to the extent that it is plausible, it strongly supports the idea that phenomenal character is not merely dispositional.

Our argument against enactivism is simply this, then: enactivism entails that phenomenal consciousness is a dispositional property; but phenomenal consciousness is not a dispositional property; therefore, enactivism is false.³²

It might be objected that the first premise is based on a superficial interpretation of enactivism. On one interpretation, the enactivist view is not that the phenomenology *consists in* dispositions, but rather that the phenomenology is *as of* dispositions. That is, the phenomenology itself is an occurrent, manifest property, it is just that what is phenomenologically manifested is a feeling of tendency and potentiality.

This is an intriguing interpretation of the enactivist approach, and its result is certainly an interesting and quite insightful take on the facts of phenomenology.³³ However, in this version there is nothing about the enactivist view to suggest that the vehicles of phenomenal consciousness are extended. It is perfectly possible for a brain in

a vat to duplicate perfectly one's phenomenology as of dispositions, tendencies, and potentialities. More generally, so interpreted the enactive approach guarantees no special connection between phenomenal consciousness and action: a Strawsonian weather-watcher could have the exact same phenomenology the enactivist claims we have.³⁴ In other words, in this interpretation the enactive aspect of phenomenology pertains to its content, and has no implication for its vehicles.³⁵

The overall argument against the objection from enactivism may thus be cast as a dilemma. Either enactivism construes the phenomenology as consisting *in* dispositions, or it construes it as a phenomenology *of* dispositions. If the former, then enactivism is inconsistent with the fact that phenomenal consciousness is a non-dispositional property. If the latter, then enactivism does not support phenomenal vehicle externalism or an "extended consciousness" view. Either way, it poses no serious threat to the phenomenal vehicle internalism presupposed by our strongly internalist, Cartesian picture of mind.

It may still be, of course, that phenomenal vehicle externalism ought to be accepted for some other reason. But if so, a compelling reason needs to be offered. On the face of it, there is no reason to think that the vehicles of phenomenal consciousness are extended. Indeed, everything known about the tokening of conscious states from work on the physical correlates of consciousness suggests that consciousness has a *neural* correlate, i.e., is tokened inside the nervous system (and hence inside the head). Ultimately, the question is empirical, and we do not wish to rule out the possibility that some evidence may yet arise that would support phenomenal vehicle externalism. Nor do we wish to deny that the *possibility* of extended conscious experiences. But at this time every indication favors the contingent truth of phenomenal vehicle internalism.

6. Conclusion

The traditional, pre-twentieth-century picture of mentality had two dimensions. The first was that there is a list of severally necessary and conjointly sufficient conditions for mentality, the second that all these conditions are always intra-cranial. The kind of externalist approaches to mentality that have flourished in the second part of the twentieth century have held on to the first dimension but rejected the second, claiming

that some of the necessary conditions on mentality are sometimes extra-cranial. On the picture of mentality proposed here, it is the first dimension rather than the second that ought to be rejected. On the picture we recommend, for most mental states, what makes them mental is that they bear the right relationship to prototypically, paradigmatically, unquestionably and unproblematically mental states; and what makes the latter mental is that they are phenomenally intentional states, that is, states that have an intentional content purely in virtue of their phenomenal character.

Because a strongly Cartesian outlook is highly plausible (and is certainly commonly thought to be *most* plausible) for phenomenally intentional states, once one adopts this picture of mentality and mental status, what are often considered challenges to a strongly Cartesian outlook on the mind appear to dissolve. In particular, the challenge allegedly presented by the extended mind hypothesis can be entirely disarmed. We have argued for a three-part position on this challenge: (a) the extended states attributed to subjects under the extended mind hypothesis turn out to be marginal cases (if cases at all) of mentality; (b) indeed, there is no deep fact of the matter as to whether they are or are not such cases; and (c) to the extent that they are, it is only because they bear the right relationship to mental states that are purely internal in every sense.³⁶

References

- Adams, F. and K. Aizawa 2001. "The Bounds of Cognition." *Philosophical Psychology* 14: 43-64.
- Brentano, F. 1874. *Psychology from Empirical Standpoint*. Edited by O. Kraus. English edition L. L. McAlister. Translated by A. C. Rancurello, D. B. Terrell, and L. L. McAlister. London: Routledge and Kegan Paul, 1973.
- Block, N.J. 1978. "Troubles with Functionalism." *Minnesota Studies in the Philosophy of Science* 9: 261–325.
- Burge, T. 1979. "Individualism and the Mental." *Midwest Studies in Philosophy* 4: 73-122.
- Chalmers, D.J. 2002. "The Components of Content." In D. J. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford and New York: Oxford UP.
- Chalmers, D. J. 2004. "The Representational Character of Experience." In B. Leiter (ed.), *The Future of Philosophy*. Oxford: Oxford UP.

- Chalmers, D.J. 2006. "Two-Dimensional Semantics." In E. LePore and B. Smith (eds.), *Oxford Handbook of the Philosophy of Language*.
- Clark, A. and D.J. Chalmers 1998. "The Extended Mind." *Analysis* 58 : 7-19.
- Dennett, D.C. 1987. *The Intentional Stance*. Cambridge MA: MIT Press.
- Dretske, F.I. 1995. *Naturalizing the Mind*. Cambridge MA: MIT Press.
- Georgalis, N. 2006. *The Primacy of the Subjective*. Cambridge MA: MIT Press.
- Giere, R.N. 2006. *Scientific Perspectivism*. Chicago: University of Chicago Press.
- Harman, G. 1990. "The Intrinsic Quality of Experience." *Philosophical Perspectives* 4: 31-52.
- Horgan, T. and G. Graham in press. "Phenomenal Intentionality and Content Determinacy." In R. Schantz, *Prospects for Meaning*. Amsterdam: de Gruyter.
- Horgan, T. and J. Tienson 2002. "The Intentionality of Phenomenology and the Phenomenology of Intentionality." In D. J. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford and New York: Oxford UP.
- Horgan, T., J. Tienson, and G. Graham 2004. "Phenomenal Intentionality and the Brain in a Vat." In R. Schantz (ed.), *The Externalist Challenge: New Studies on Cognition and Intentionality*. Amsterdam: de Gruyter.
- Hurley, S.L. 1998. *Consciousness in Action*. Cambridge MA: Harvard UP.
- Kriegel, U. 2003a. "Is Intentionality Dependent upon Consciousness?" *Philosophical Studies* 116: 271-307.
- Kriegel, U. 2003b. "Consciousness as Sensory Quality and as Implicit Self-Awareness." *Phenomenology and the Cognitive Sciences* 2: 1-26.
- Kriegel, U. 2007a. "A Cross-Order Integration Hypothesis for the Neural Correlate of Consciousness." *Consciousness and Cognition* 16: 897-912.
- Kriegel, U. 2007b. "Intentional Inexistence and Phenomenal Intentionality." Forthcoming in *Philosophical Perspectives* 21: 307-340.
- Loar, B. 2002. "Phenomenal Intentionality as the Basis for Mental Content." In D. J. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford and New York: Oxford UP.
- McGinn, C. 1988. "Consciousness and Content." In *Proceedings of the British Academy* 76: 219-239. Reprinted in N. J. Block, O. Flanagan, and G. Güzeldere (eds.), *The Nature of Consciousness: Philosophical Debates*. Cambridge MA: MIT Press, 1997.
- McGinn, C. 2008. "Consciousness as Knowingness." This volume.
- Marr, D. 1982. *Vision*. San Francisco: WH Freeman Publishers.
- Milner, A. D., and M. A. Goodale 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.
- Noë, A. 2004. *Action in Perception*. Cambridge MA: MIT Press.
- Noë, A. 2006. "Experience of the World in Time." *Analysis* 66: 26-32.

- O'Regan, J.K. and A. Noë 2001. "What It is Like to See: A Sensorimotor Approach to Perceptual Experience." *Synthese* 129: 79-103.
- Pitt, D. 2004. "The Phenomenology of Cognition; or *What Is It Like to Think that P?*", *Philosophy and Phenomenological Research* 69: 1-36.
- Putnam, H. 1975. "The Meaning of 'Meaning'." In his *Mind, Language, and Reality*. Cambridge: Cambridge UP.
- Rorty, R. 1970. "Incorrigibility as the Mark of the Mental." *Journal of Philosophy* 67: 399-429.
- Rosch, E.H. 1973. "Natural categories." *Cognitive Psychology* 4: 328-350.
- Searle, J.R. 1991. "Consciousness, Unconsciousness, and Intentionality." *Philosophical Issues* 1: 45-66.
- Searle, J.R. 1992. *The Rediscovery of Mind*. Cambridge MA: MIT Press.
- Siewert, C.P. 1998. *The Significance of Consciousness*. Princeton NJ: Princeton UP.
- Stich, S. 1979. "Autonomous Psychology and the Belief-Desire Thesis." *The Monist* 61: 571-91.
- Strawson, G. 1994. *Mental Reality*. Cambridge MA: MIT Press.
- Strawson, G. 2004. "Real Intentionality." *Phenomenology and the Cognitive Sciences* 3: 287-313.
- Tartaglia, J. 2008. "Consciousness, Intentionality, and the Mark of the Mental: Rorty's Challenge." Ms.
- Williford, K. 2005. "The intentionality of consciousness and the consciousness of intentionality." In G. Forrai and G. Kampis (eds.), *Intentionality: Past and Future*. Dordrecht: Kluwer Academic Publishers.
- Yablo, S. 1997. "Wide Causation." *Philosophical Perspectives* 11: 251-81.

¹ See Horgan and Tienson 2002, Horgan, Tienson, and Graham 2004, Kriegel 2003a, Kriegel 2007b.

² These conditions are summarized by Clark and Chalmers (1998: 17) as follows: "First, the notebook is a constant in Otto's life—in cases where the information in the notebook would be relevant, he will rarely take action without consulting it. Second, the information in the notebook is directly available without difficulty. Third, upon retrieving information from the notebook he automatically endorses it. Fourth, the information in the notebook has been consciously endorsed at some point in the past, and indeed is there as a consequence of this endorsement."

³ The story is fleshed out much more in Clark and Chalmers 1998, in a way that makes it sound much more plausible than it might strike the reader from our presentation. We rest content with this emaciated presentation because we do not intend to contest EMH on the grounds that the leading thought-experiment in its favor fails to support it.

⁴ This leaves open the possibility of denying psychological externalism, not by rejecting the familiar arguments in its favor (Putnam 1975, Burge 1979), but by denying the content-based individuation of mental states—or at least of mental states with externalistically individuating content (Stich 1979). Content externalism is only an existential, not universal claim, so a defender of psychological externalism could resort to a kind of divide and conquer strategy whereby mental states with internalistically individuating

contents individuate by their contents and mental states with externalistically individuating content individuate by, say, syntax or functional role.

⁵ This may well be Chalmers' view, given his defense of content internalism in Chalmers 2002, 2006, and other places.

⁶ These ideas are pursued in the writings of, among others, Nicholas Georgalis (2006), Brian Loar (2002), Colin McGinn (1988, 2008), John Searle (1991, 1992), Charles Siewert (1998), Galen Strawson (1994, 2004), and of course ourselves (Horgan and Tienson 2002, Horgan, Tienson, and Graham 2004, Kriegel 2003a, Kriegel 2007b). It is unclear to us that anybody (else) has committed to all of these ideas in print, and some have positively rejected one or two. But all of the above are committed to the loose cluster of these theses in the sense that they endorse *most* of them.

⁷ This is so inasmuch as coherence and explanatory unity are central theoretical virtues in light of which to assess overall theories (and "outlooks").

⁸ On this, see especially Horgan and Tienson 2002, Loar 2002, McGinn 1988, and Siewert 1998.

⁹ In Horgan and Tienson 2002, and in Horgan, Tienson and Graham 2004, a distinction is drawn between phenomenal intentionality and *externalistic* intentionality. It is claimed that some thought constituents, e.g., those expressible linguistically via proper names or natural-kind terms, have the following two features. First, such thought-constituents have *externalistic referential purport* that is determined wholly phenomenologically: there are phenomenologically determined reference-eligibility conditions that require the world to be a certain way in order for the given thought-constituent to successfully refer. Such externalistic referential purport is an aspect of phenomenal intentionality. Second, when a singular or natural-kind thought-constituent actually *refers* to some item in the experiencer's environment (viz., some object or some natural kind), this is because that item uniquely satisfies the pertinent reference-eligibility conditions. So full-fledged reference, for such thought-constituents, is a form of intentionality that depends constitutively not just on phenomenology but also on the experiencer's wider environment. Because of the constitutive environment-involving aspects, reference (for such thought-constituents) is termed *externalistic* intentionality.

¹⁰ See Georgalis 2006, Loar 2002, Horgan and Tienson 2002, Horgan, Tienson, and Graham 2004, Kriegel 2007b, Searle 1992, Strawson 2004, as well as Chalmers 2004; but see also McGinn 1988, who embraces externalism about phenomenally intentional content.

¹¹ A more sustained argument on behalf of cognitive phenomenology is provided by Pitt (2004). See also Kriegel 2003b, Siewert 1998, and Strawson 2004.

¹² Here see mainly Loar 2002, Kriegel 2003a, McGinn 1988, and Searle 1991, 1992. A particularly strong version of this claim would be that only phenomenally conscious states have intentionality, genuine intentionality, at all; this claim is made in Georgalis 2006 and Strawson 2004.

¹³ Some conscious states too might well have a derivative form of intentionality, over and above their phenomenal intentionality. Consider conscious thoughts deploying singular thought-constituents purporting to refer to concrete individuals, and/or predicative thought-constituents purporting to refer to natural kinds. According to Horgan and Tienson 2002 and Horgan, Tienson, and Graham 2004 (cf. note 6), the phenomenal intentionality of such thoughts fixes what could count as eligible referent-items, for the singular and natural-kind thought-constituents. Actual *reference* for such thought-constituents, on the other hand, is a derivative form of intentionality that Horgan et. al. call externalistic intentionality. It is derivative in the sense that it depends in part upon phenomenal intentionality (while also depending in part on the presence in the experiencer's environment of items that satisfy the phenomenally determined reference-eligibility conditions).

¹⁴ See also Horgan and Graham, in press.

¹⁵ A different kind of argument can be found in Strawson (2004), who offers an argument from pansematicism: by the lights of causal-covariational accounts of intentionality, just about anything will turn out to have intentionality not derivatively.

¹⁶ Thus while it is difficult to imagine a cognitive system that has just one belief – something that at some point inspired many philosophers to recoil from atomism about the propositional attitudes – there is no difficulty in imagining a creature that comes into being, has a single momentary conscious experience, and then disappears.

¹⁷ At any rate, it is natural to construe Dennett as a fictionalist about intentional mental states, and many have so construed him—even though his writings are persistently somewhat equivocal on the matter.

¹⁸ Not everything that is useful must be a fiction. Thus, it is useful to regard cars as instruments of transportation, but that is not a fiction: they really are instruments of transportation. Our view of instrumentally ascribed mental states is similar.

¹⁹ It might be objected that talk of *inferential* connections between non-phenomenal states and phenomenally intentional states makes no clear sense unless the former are construed as having pre-existing intentionality that grounds the notion of an “inferential connection.” But on the picture we mean to be proposing, that worry does not arise. An inferential connection, in the relevant sense, is a typical-cause relation whose obtaining contributes to the overall hermeneutic construal of the agent as an “intentional system.” Furthermore, although worries about radical content-indeterminacy might arise if the only constraint on content-assignments were that they should collectively constitute a Davidson-style “radical interpretation” of the agent’s internal states as mental states, bodily motions as purposive actions, and verbal outputs as assertions, nevertheless on our picture such radical indeterminacy is prevented by the role of phenomenally conscious states as content-determinate “anchor points” in an overall-acceptable radical interpretation. (For elaboration of this claim about content determinacy, see Horgan and Graham in press.)

²⁰ Such causal integration will also underwrite content-determinacy for the non-conscious states, via their inferential-connection relations to phenomenally intentional states, whose determinate intentionality is intrinsic to their phenomenal character, for those philosophers who take intrinsic content-determinacy to be the key difference between phenomenally intentional states and other intentional states (Horgan and Tienson 2002, Horgan and Graham in press).

²¹ See Tartaglia 2008 for a much more thorough discussion of this point.

²² It is an empirical matter, of course, which instances are prototypical. As it happens, studies show that the sparrow is the leading folk prototype of a bird.

²³ It will be useful to distinguish two kinds of prototypicality, weak and strong. Some birds (e.g., sparrows) are maximally prototypical instances of the category *bird*, whereas others (e.g., bluejays) are sufficiently close to maximal prototypicality that they are aptly considered prototypical birds themselves (although not *maximally* prototypical birds). The sparrows are *strongly* prototypical birds, whereas the bluejays are *weakly* prototypical birds. In terms of this distinction, the double proposition to which the present note is appended pertains to weak prototypicality for the category *mental state*. Some phenomenally intentional states are more prototypical than others, and we think it is empirically plausible that the *strong* prototypes have to do with visual experiences and occurrent thoughts (which we take to have a cognitive phenomenology). Visual experiences are the mental states where there is perhaps the clearest and most undeniable phenomenology and intentionality, as well as a close connection between them: the phenomenology seems most acutely to constitute the intentionality there. This is perhaps why claims of the transparency of experience (Harman 1990) are most immediately compelling when applied to visual experiences. It is also clear that, once one accepts the idea of cognitive phenomenology, conscious occurrent thoughts seem to be transparent in the same way: we can be introspectively aware of nothing about them except their contents.

²⁴ It might be objected that the folk concept of mentality, or of mental state, cannot be and/or ought not to be a prototype concept, because it is a natural kind concept, and so must appeal to necessary and sufficient underlying nature. It seems to us, however, that there is no real tension between being a natural kind concept and being a prototype concept. A natural kind prototype concept would be one for which the relevant relationship non-prototypical instances would have to bear to prototypical ones is that of (probably exact) similarity with respect to underlying nature. For an unnatural-kind prototype concept, the relationship would have to be one of similarity with respect to manifest features.

²⁵ Brentano's texts also are most naturally construed, we would maintain, as committed to the view that genuine consciousness is what is nowadays often called *phenomenal* consciousness; Brentano did not acknowledge any states that are conscious-as-opposed-to-unconscious but fail to be phenomenally conscious.

²⁶ This is a problem because, obviously, not only phenomenally conscious states are mental states. We mentioned the main arguments for taking only phenomenal states to have non-derivative intentionality in the previous section.

²⁷ Clark and Chalmers (1998) concede that conscious experiences are never extended (as we will see in §5). It seems to follow that extended states are not potentially conscious. So if they are to count as mental states, the Searlean mark thesis must be rejected.

²⁸ Although someone might say, not implausibly, that skin and skull just happen to be where certain sensory transducers are located, and that there is something special about those. We will not pursue this line of thought here.

²⁹ For present purposes, we can remain neutral on whether the tokens of phenomenally intentional state-types are neural state-tokens, or instead are distinct state-tokens that bear some intimate non-identity relation (e.g., *being constituted by*) to neural state-tokens. So, for present purposes, the claim that phenomenally intentional states (state-types) are *internally tokened* is also neutral between these two options.

³⁰ Envatted-brain scenarios usually involve a systematic coupling between the brain, on one hand, and on the other hand the external device (usually envisioned as a supercomputer) that monitors the brain's motor-output signals and feeds it sensory-input signals. Arguably, the states of the external device could be regarded as aspects of an extended mind that is physically realized by the coupled brain-computer system as a whole. (Or at any rate, a coupled brain/computer system that is a phenomenal duplicate of someone like Otto would be, arguably, just as good a *candidate* for having an extended mind as is Otto himself.) But this need not affect our main point in citing the envatted-brain example, because the example can be tailored to avoid the point. Let the pertinent scenario be an envatted brain that receives randomly-generated inputs from a long-term surrounding electrical storm that just happens, via cosmic coincidence, to give the brain an ongoing phenomenal mental life that exactly matches yours.

³¹ In fact, the contrary is more plausible: Clark and Chalmers' appeal to the parity principle, without which the case of Otto and Inga would not support the EM hypothesis, seems to be an implicit appeal to something like a mark-of-the-mental thesis.

³² Even if one were to embrace the implausible view that phenomenal consciousness is a dispositional property, a somewhat less implausible version of this generically implausible position would identify phenomenal consciousness not with the dispositions of an embodied agent to interact with objects of its environment, but rather with the dispositions of a cognitive control-system to dynamically generate motor-output commands in response to ongoing sensory representations that dynamically arise "inside the transducers." One advantage of this latter version of dispositionalism is the verdict it yields concerning a recently envatted brain that was embodied until yesterday but that now is receiving nonveridical sensory inputs in response to its motor-control outputs. Whereas the version of dispositionalism we have been

criticizing evidently entails that this recently envatted brain has no phenomenal consciousness, the alternative form of dispositionalism at least accords the brain the phenomenal consciousness that it surely possesses. (Likewise, *mutatis mutandis*, for someone who has become totally paralyzed.)

³³ At the same time, some implausibility attaches to the idea that *all* phenomenology is dispositional in this sense. Even if the dispositional account is right story to tell about the way in which the hidden parts of the apple are phenomenologically manifest, it seems a little strained to say that the same story is appropriate for the part of the apple that is in plain view. The way in which the surface of the part of the apple facing us is present in phenomenology seems not to require any potentialities or tendencies. If so, it cannot be that phenomenology is *entirely* dispositional in this sense. Which also seems to suggest that, even to the extent that our actual phenomenology *is* dispositional, that cannot be a constitutive aspect of phenomenology, something without which there cannot be any phenomenology.

³⁴ Strawson (1994) describes a race of weather-watchers, creatures who cannot move but have a rich conscious life pertaining to the weather around them.

³⁵ Of course, insofar as the content of phenomenology (especially perceptual-experiential phenomenology) is as of dispositions, these are dispositions of one's apparent body to interact in various ways with apparent objects in one's apparent environment—not dispositions of one's cognitive-control system to generate ongoing motor-output commands in response to one's ongoing sensory representations. But it would be a serious fallacy to conclude, on the basis of this fact about the *content* of perceptual phenomenology, that the *vehicles* of perceptual phenomenology extend beyond the skull and skin. Susceptibility to that very fallacy, it appears to us, largely explains the attractions that some people feel to the extended mind outlook.

³⁶ For useful conversations, we would like to thank Yali Corea-Levy, Brian Fiala, Ron Giere, Kristie Miller, Cole Mitchell, Shaun Nichols, and Anne Steadman.