

Stop Caring about Consciousness

Peter Carruthers
University of Maryland

ABSTRACT. The best empirically grounded theory of first-personal phenomenal consciousness is global workspace theory. This, combined with the success of the phenomenal-concept strategy, means that consciousness can be fully reductively explained in terms of globally broadcast representational content. So there are no qualia (and there is no mental paint). As a result, the question of which other creatures besides ourselves are phenomenally conscious is of no importance, and doesn't admit of a factual answer in most cases. What is real, and what does matter, is a multi-dimensional similarity space of functionally organized minds.

1. INTRODUCTION

Our topic is phenomenal consciousness. This is the kind of consciousness that gives rise to “hard”-problem thought experiments, such as the conceivability of zombies and the explanatory gap. Phenomenal consciousness is fundamentally first-personal. Talk of the “feel” of experience and of what experiences are “like” are really just invitations to one's hearers to pay attention to their own experiences to verify what one is saying from their own introspection (Block 1995). It is these introspection-based first-person thoughts that define the subject matter. We can, in addition, formulate a public concept of phenomenal consciousness, of course. But this is best characterized as: whatever property it is that gives rise to

the “hard”-problem thought experiments (Carruthers and Veillet 2017). That is what I will assume here.

There are numerous other ways in which the term “consciousness” is used, of course, and these need to be sharply distinguished from our target (Rosenthal 2005). Indeed, the field of consciousness studies is littered with confusion because people fail to keep sight of these important distinctions, slipping illegitimately from ascribing consciousness in one sense to consciousness in another. A couple of points are worth mentioning briefly here.

First of all, phenomenal consciousness is a variety of mental-state consciousness. It is the mental states of a person that are *like* something to undergo, and that have a distinctive subjective *feel*. People/agents are phenomenally conscious derivatively, by virtue of undergoing phenomenally conscious mental states. Mental states can also be conscious in another sense, however, by being accessible to inform reasoning and decision making, as well as giving rise to long-term memories and issuing in verbal reports. This is *access* consciousness, which is third-personal (functionally defined), and is at least conceptually distinct from phenomenal consciousness (Block 1995). As we will see shortly, however, access consciousness, in its contemporary guise of global workspace theory, provides the best candidate for the neural/functional correlate of phenomenal consciousness, at least.

People (and other animals) can also be said to be conscious in a quite different sense. This is *creature* consciousness, which admits of both transitive and intransitive varieties. Intransitive creature consciousness is a matter of being awake rather than asleep, or conscious rather than comatose. Almost all living creatures undergo sleep/awake cycles, of course, and so are sometimes conscious in this sense. Transitive creature consciousness, on the other hand, coincides with perceptual consciousness. A creature can be conscious *of* some aspect of the world or of its own body, meaning that it perceives it. Thus, a cat that sees a mouse scuttling along the wall can be said to be conscious of the movement of the mouse. Both forms of creature consciousness, like access consciousness, are third-personal and functionally defined. And likewise, both admit of degrees (again, also like access consciousness). A creature can be more or less awake (one can be partially conscious), and perceptual consciousness can be more or less rich in content, both within an individual and across different species of animal.¹

There is much to learn about perception in ourselves and other creatures, of course, as well as about the varieties and richness of perceptual consciousness across species. Likewise, there is much yet to be learned about the nature and

1. In Carruthers (2019) I place quite a bit of weight on the fact that all of the functionally defined forms of consciousness admit of degrees across species, whereas first-person phenomenal consciousness is all-or-nothing—any particular mental state in oneself is either *like something* to undergo or it isn't. This plays an important role in what I call “the negative semantic argument” for the conclusion also defended in the present article; but I won't need to rely on it in the discussion that follows.

function of sleep and dreaming, as well as about the mechanisms that govern waking and alertness. These are hard scientific questions, well worth studying. But they raise no problems for a scientific-naturalist, basically physicalist, view of the world, any more than mentality itself does. Third-personal forms of consciousness are straightforwardly amenable to scientific study and, ultimately, scientific explanation.

Phenomenal consciousness, in contrast, is thought by many to give rise to deep problems for physicalist forms of naturalism. This is because one can see that no matter how much one learns about the physical, functional, and representational properties of one's own phenomenally conscious states, one will still face an explanatory gap when seeking to account for those states' phenomenal properties. For one will be able to think, "all of that [physical/functional/representational] stuff might have been true, and yet *this feel* might have been different or absent." Likewise, there seems to be no difficulty in conceiving of a creature that is like oneself in all physical, functional, and representational respects but who nevertheless lacks *these* kinds of (phenomenally conscious) mental states altogether. That is, one can conceive of a zombie version of oneself. In consequence, many have wanted to conclude that phenomenal consciousness involves a distinctive set of nonphysical properties (qualia) that are responsible for the felt qualities of one's experience (most notably, Chalmers 1996).

I will argue here, however, that the so-called "hard" problem of consciousness results entirely (and, in a sense, trivially) from the differing perspectives one can take on one-and-the-same experiential state (third person versus first person). No extra properties enter the world with phenomenal consciousness, and nothing mysterious is implicated. Moreover, when one attempts to project one's first-person conception of phenomenal consciousness across species, the result is indeterminacy, and there are no real facts of the matter waiting to be discovered.

2. GLOBAL WORKSPACE THEORY

Over the past 40 years, numerous labs around the world have joined the hunt for the neural correlates of (phenomenal) consciousness, using a variety of scientific methods and technologies. The result has been a proliferation of theories, including global workspace theory (Baars 1988; Dehaene 2014), integrated information theory (Tononi 2008; Tononi and Koch 2015), higher-order thought theory (Lau and Rosenthal 2011; Rosenthal 2005), and a theory that locates phenomenal consciousness in the contents of a form of fragile short-term memory (Block 2007, 2011). I have examined these (and other) theories in some detail elsewhere (Carruthers 2019). Here I will just sketch some of the evidence that supports global workspace theory, as well as some reasons for rejecting Block's suggestion that consciousness is correlated with a richer form of short-term memory, since this view seems especially popular among philosophers.

Particularly important in the investigation of the neural correlates of consciousness have been experimental paradigms in which the stimuli can be kept exactly the same, but where consciousness can be either present or absent. For example, in backward-masking experiments, very briefly presented stimuli are followed immediately by a follow-up stimulus, which often has the effect of rendering the first image invisible. If the timing and intensity of the stimuli are carefully titrated, one can achieve a situation in which one-and-the-same type of stimulus is consciously perceived on about 50 percent of trials. One can then use fMRI, MEG, and/or EEG to investigate what happens differently in the brain when the same stimuli are experienced consciously as opposed to unconsciously.²

The general finding in such experiments is that in cases where a stimulus remains unconscious (albeit being successfully processed up the level of semantic meaning, perhaps, as well as priming subsequent responses), there is reverberating activity in mid- and high-level sensory cortices at the “back” of the brain that rapidly fades away. When the stimulus is consciously experienced, in contrast, there is coordinated activity involving prefrontal, parietal, and temporal cortices as well. Multiple forms of inquiry suggest that attention plays an important role in this process. Attentional signals boost the activity of target populations of neurons while suppressing competitors, generally resulting in a sort of global ignition that enters the targeted contents into working memory, thereby rendering them widely available to a range of different systems capable of remembering, describing, or otherwise responding to them. In consequence, many theorists now identify consciousness with the contents of working memory (or the contents of the “global workspace”; Carruthers 2019; Dehaene 2014; Mashour et al. 2020).³

Critics have objected that the prefrontal activity one observes when stimuli are consciously experienced might not be a correlate of consciousness itself, but rather the engagement of sustained attention and/or response preparation (Michel 2017). A variety of clever experiments make this interpretation unlikely, however. One of these required participants to make a response on every trial, whether the stimulus was consciously experienced or not (Salti et al. 2015). Since participants often responded correctly at above-chance levels even when the stimuli were unseen, one can assume that a form of blindsight-like action-facilitation was at work. The experimenters then trained pattern classifiers on the resulting brain-imaging data to see

-
2. fMRI (functional magnetic resonance imaging) has quite good spatial but poor temporal resolution. EEG (electroencephalogram) has excellent temporal resolution but poor spatial resolution. MEG (magnetoencephalography) has intermediate virtues and drawbacks, both spatial and temporal. Sometimes more than one of these methods can be combined together in a single experiment.
 3. Or rather, they identify consciousness with actively maintained forms of working memory. For there is recent evidence that information can be passively retained, and readily recovered, for short periods of time in the form of temporarily altered synaptic weights (Lewis-Peacock et al. 2012; Trübtschek et al. 2018). I myself think that this shouldn't really be counted as a form of working memory at all. Rather, it belongs alongside the phenomena that motivated people to talk about “long-term working memory” (Ericsson and Kintsch 1995).

where in the brain the content of the stimulus was represented when consciously seen, contrasting this with cases where participants responded correctly but without undergoing conscious experience. In the former case, those contents reached much deeper into prefrontal regions while being processed more robustly there.

One particular style of experiment is especially telling. This makes use of binocular rivalry, and has been conducted with monkeys as well as humans. In binocular rivalry, distinct images are simultaneously presented to the two eyes. At any one time, only one of the images is consciously experienced, with experience flipping from one image to another at irregular intervals because of neural adaptation. Given the right experimental conditions and/or as a result of pretesting, one can sometimes reliably predict which stimulus will be experienced within a specific time frame. In two recent experiments conducted with monkeys, the experimenters recorded from populations of neurons in higher-level visual areas, and also in regions of lateral prefrontal cortex, after first discovering which populations represented which of the two stimuli. The finding was that prefrontal content-specific neurons are active in purely passive viewing conditions when and only when the corresponding stimulus is dominating (Kapoor et al. 2020; Panagiotaropoulos et al. 2012).^{4, 5}

Block (2007, 2011) concedes the evidence that access consciousness involves prefrontal representation, while arguing that phenomenal consciousness is nevertheless something distinct: accessible for report when targeted by attention, but with richer contents than can enter working memory at any one time. For when a rich stimulus array of letters or shapes is briefly presented on a screen, participants report conscious experience of the entire array, despite only being able to identify about four items (the standard limit for the contents of working memory). Yet, if the location of any one target is retrospectively cued, participants can report it, suggesting that their conscious experience was, indeed, rich, overflowing what can simultaneously enter working memory (Landman et al. 2003; Sligte et al. 2008;

-
4. One might worry that questions about phenomenal consciousness in animals are begged by appealing to experiments conducted with monkeys. But in fact, the experiments just presuppose a (limited) form of access consciousness in monkeys, and their evidential role is to demonstrate that the prefrontal activity observed when humans undergo phenomenally conscious states is not due to either response preparation or the engagement of attention. For attention is engaged whichever of the two stimuli dominates in binocular rivalry; and the monkeys remained entirely passive: no response was required of them.
 5. Block (2019) objects against this paradigm that, while it rules out response-preparation as an explanation of the prefrontal activity that correlates with conscious experience, it can't rule out *thoughts about* the conscious experiences in question. Perhaps content-specific prefrontal activity reflects, not the conscious status of the percepts in question, but rather down-stream thoughts about those percepts. This is possible, but unlikely given the data. Block's suggestion cannot explain why participants are *always* (or almost always) thinking about the currently conscious stimulus on a trial-by-trial basis, nor why the onset and offset times of their thoughts should so closely coincide with switches in conscious status. What one thinks about is not deterministically controlled by what one sees, surely. Moreover, we know from direct-intervention studies conducted with humans that stimulation of prefrontal cortex can be sufficient to induce, not just thoughts, but conscious visual hallucinations (Blanke et al. 2000; Vignat et al. 2000). This, too, supports a critical role for prefrontal cortex in conscious experience.

Sperling 1960). Block thus identifies phenomenal consciousness with the rich contents of a fragile short-term memory system located in mid- or higher-level sensory cortices.

Critics have acknowledged the existence of fragile short-term sensory memory, and its richness; but they have denied that the data demonstrate that its contents are conscious. Rather, the rich stimulus array can be represented unconsciously until a cue targets a particular item with attention, resulting in it *becoming* conscious (Cohen and Dennett 2011; Cohen et al. 2016). People's impression of richness when experiencing the initial stimulus can be explained by a combination of consciously experienced *gist* content (e.g. "a bunch of letters arranged in rows") together with conscious experience of just a few items (Cohen et al. 2011; Wu and Wolfe 2018) no doubt reinforced by the background folk-psychological assumption of the richness of experience in general, resulting from ignorance of the frequent saccades made by one's eyes, which constantly sample the environment about three times every second.

The data Block appeals to in his support are not persuasive, then. In addition, one reason for positively rejecting Block's claims is that other retro-cuing experiments have found that a cue that is presented following stimulus offset can improve even conscious awareness of just a single faint stimulus (Sergent et al. 2013; Thibault et al. 2016; Xia et al. 2016). If the contents of fragile short-term memory are always conscious, then a retro-cue should be ineffective: the stimulus should already be conscious. And if it is replied that the stimulus may give rise to a representation in short-term memory that is too faint to be conscious in the absence of directed (cued) attention, then that undermines the probative value of the original Sperling-style experiments. For it can be said that the contents of fragile short-term memory, in general, remain unconscious until targeted by attention.

Another set of experiments causing problems for any view that identifies consciousness with the relatively rich contents of fragile short-term memory is provided by Tsubomi et al. (2013). They first tested working memory limits for colored squares in a standard way, by presenting an array of them, followed by a retention interval, followed by a single square selected from one of the original positions. But now one half of that square retained the original color, whereas the color of the other half was drawn from elsewhere in the initial array. Participants just had to say which was the original color of the square. Tsubomi and colleagues found the standard four-item limit overall. But then they re-ran the test without any retention interval: at the very moment when the other squares disappeared, the remaining one became color divided, with one half retaining the original color and the other color drawn from elsewhere in the initial array. There was no difference in performance: it still leveled off when more than four items overall were presented (and by just the same amount).

This finding would be mysterious if it were true, as Block claims, that an entire array of up to around eight or more items can be consciously experienced at the same time. For if the item that remains when the others disappear were already

consciously experienced, then one would think one would immediately know which side of it was the original color and which side had just changed. At any rate one would predict that performance would be *better* in the immediate-change case than when there is a retention interval of nearly a second. But in fact, there was no difference, suggesting that the limits on conscious experience and the limits on working memory are the same.^{6, 7}

In what follows, then, for the reasons briefly sketched here (together with others; see Carruthers 2019), I propose to take for granted that phenomenal consciousness lines up one-for-one with the contents of the global workspace (with access consciousness, in other words). I also propose to assume (but won't argue here) that the contents in question are always representational. That is to say, there are no differences in phenomenal consciousness that aren't accompanied by representational differences; and likewise, when representations figure within the global workspace, there are no representational differences that don't impact phenomenology. So there is no "mental paint" (Block 2010), and the phenomenal properties of things like moods and the painfulness of pain can be captured in representational terms. This has been argued elsewhere by myself and others (Carruthers 2018, 2019; Cutter and Tye 2011; Tye 1995, 2000). I don't propose to revisit those arguments here.⁸

3. THE PHENOMENAL-CONCEPT STRATEGY

Even if phenomenal consciousness correlates one-to-one with globally broadcast content, it doesn't follow that phenomenal consciousness *is* globally broadcast content, of course. But that is the view I will sketch arguments for here. I will

6. I have simplified my presentation of these findings for present purposes. In fact, participants kept their eyes on a central fixation point, and the squares were presented on either side of it. *Two* squares then remained on the screen, one in each hemi-field. The finding of an overall four-item limit results from a two-item limit within each hemisphere of the brain, which is also reflected in a distinctive neural signature (the contralateral delay activity, or CDA), which increases with working memory load but levels off at the two-item limit. Notice that these findings are consistent with, but do not require, a "slot-based" account of working-memory limits (Adam et al. 2017). They are equally consistent with resource-based accounts (Ma et al. 2014).

7. Why is it that these experiments display a four-item limit, whereas Sperling-style retro-cue experiments seemingly extend that limit? The difference is likely to be that in the present experiments, the stimulus-to-be reported over-writes the original. In contrast, in Sperling-style ones the to-be-reported location is indicated indirectly, thus not interfering with the fragile short-term memory that can then be targeted by attention.

8. I should note that whenever I talk about *content* in this paper I intend (unless otherwise indicated), specifically to refer to *nonconceptual* (analog, fine-grained) content. This is because I believe that only nonconceptual content makes a constitutive contribution to the phenomenal properties of our experience, giving rise to hard-problem-type thought experiments. See Carruthers and Veillet (2011, 2017). But nothing substantive will turn on this assumption for our purposes here.

suggest that all of the facts surrounding phenomenal consciousness (including the “hard”-problem thought experiments) can be fully reductively explained in third-person psychological terms, allowing us to conclude that there is at least a local identity between globally broadcast content and phenomenal consciousness in humans.

The basic explanatory strategy, here, is long-standing and familiar (Balog 1999, 2012; Carruthers 2000; Carruthers and Veillet 2007; Loar 1990; Papineau 2002; Prinz 2012; Tye 1995, 2000). It is to argue that the so-called “hard” problems (the explanatory gap, the conceivability of zombies, and the rest) don’t result from any special properties attaching to our conscious experiences. (That is, they don’t result from the existence of nonphysical qualia, nor from the presence of any essentially first-personal properties of experience.) Rather, those problems arise from the special way in which we sometimes *think about* our own globally broadcast contents. Specifically, we can refer to our own access-conscious states in a way that doesn’t require explicit deployment of any of our third-person physical, functional, or representational concepts. Nevertheless, what we thereby refer to is just globally broadcast content.

The account of our first-person phenomenal concepts that I now favor follows Prinz (2012) in thinking of them as a kind of acquaintance-based indexical (Carruthers 2019). But there is nothing mysterious about the notion of *acquaintance* employed here. Suppose, for example, that one is looking at something red, in good lighting conditions. The visual system responds by producing the perceptual content, *red*. If that content is globally broadcast, then it can be used in a variety of different ways. It can be used to form first-order thoughts like, “*That* is a lovely shade of red,” where the word “that” refers to the color of the worldly object. But the very same content can also be used to form higher-order thoughts about the experience itself, like, “*This* could have been reliably caused by the presence of something green rather than red” or, “There could be a creature that is like me in all respects except that it lacks *this*” where here the word “this” refers to the content of the perceptual state itself. In both cases the content in question is received as input by the conceptual and planning systems that generate structured thoughts. The only difference is that the first-order thought depends on tacit deployment of the concept color, whereas the higher-order ones rely on tacit use of the concept perception. In both cases we have an indexical thought based in acquaintance. It is just that the way each indexical is based in acquaintance is different, depending on the concepts tacitly deployed in the background.⁹

Although phenomenal concepts, when *used* in the first person, lack conceptual connections with physical, functional, and representational concepts (enabling

9. Notice that it follows from this that phenomenal concepts aren’t *completely* disconnected from other mentalistic concepts. See Carruthers (2017) for discussion of how merely tacit deployment of concepts like perception can still leave us able to think thoughts like, “*This* [percept of red] might not have had the functional role of perception.”

us to conceive of zombies and the rest), those same concepts can also be given a distinctive third-personal characterization. Indeed, I did so in the previous paragraph, saying that they are indexical concepts used to designate one's own globally broadcast perceptual states through tacit use of the concept perception (or something similar). But why should we suppose that this is so? Indeed, doesn't it beg the question against the existence of a distinctive set of first-personal properties of experience to say that our phenomenal concepts refer, not to them, but just to globally broadcast content? What reason do believers in qualia or mental paint have for accepting such a claim?

One argument builds on considerations of explanatory simplicity. By supposing that phenomenal concepts are just acquaintance-based indexicals referring to globally broadcast contents we can explain how the various "hard"-problem thought experiments arise. Yet the explanation in question just appeals to properties that everyone agrees to exist (global broadcasting, representational content, indexical concepts, higher-order concepts). There is simply no need to postulate any additional properties. Indeed, the simplest explanation is that the explanatory gap is an illusion produced by a disconnect between our third-person physical and psychological concepts, on the one hand, and our first-personal higher-order indexical ones, on the other. Put differently, there is an explanatory gap at the level of thought (some *thoughts* about our globally broadcast states aren't entailed by third-personal facts) rather than at the level of worldly properties. The only real properties that exist are ones that are physical, functional, and representational. There are no additional first-person-individuated properties.

Here is a somewhat more direct argument for the same conclusion. When using phenomenal concepts in the first person one can, of course, conceive of the existence of a zombie: someone who is like oneself in all physical, functional, and representational respects but who nevertheless lacks anything like *this* (globally broadcast content). But given this characterization of a zombie, together with the third-personal account of phenomenal concepts sketched above, one can see in advance that the zombie, too, will be able to conceive of a zombie version of itself. Since the higher-order indexicals referring to its own globally broadcast contents lack connections with its physical, functional, and representational concepts, the zombie, too, will be able to think, "There could be a creature exactly like myself in all physical, functional, and representational respects who nevertheless lacks anything like *this*" and, "No matter how much I know about the physical, functional, and representational facts involved in perception, that still won't explain why those facts give rise to states like *this*." Since my zombie twin, too, can easily become puzzled by the same "hard"-problem thought experiments as myself, we don't need to appeal to anything other than phenomenal concepts and broadcast contents in order to explain how those thought experiments arise.

The success of the phenomenal-concept strategy hinges on the existence of phenomenal concepts, of course. But both Ball (2009) and Tye (2009) deny that there are any such things, on very similar grounds. However, as Veillet (2012) shows,

and as Carruthers (2019) also argues, their arguments lose sight of the fact that phenomenal concepts are grounded in their first-person uses, and that much talk about phenomenal consciousness doesn't really involve phenomenal concepts at all. We can, of course, entertain thoughts about phenomenal states that we have never experienced (think of Mary in her black-and-white room wondering what it is like to see red), and of course we can form concepts that defer to the phenomenal concepts of other people. But in doing so we aren't forming first-person acquaintance-based higher-order indexical concepts, of the sort that give rise to the "hard"-problem thought experiments.

4. THE QUESTION OF OTHERS' CONSCIOUSNESS

Recall the distinction we drew at the outset between the first-personal phenomenal concepts that define our subject-matter, and the third-person concept we can use to operationalize phenomenal consciousness for purposes of public discourse. The former are concepts like *THIS* (introspectively identified content of experience). The latter can be defined as "whatever property explains the hard-problem thought experiments." We now know that the latter property is globally broadcast content (or so I have suggested). So if we ask about the distribution of phenomenal consciousness across the animal kingdom using the public concept, then the answer is straightforward. It is that animals will possess the property in question (globally broadcast content) to whatever degree they share the same global workspace architecture as ourselves. And as we will see later, what exists in nature is a complex cross-cutting multidimensional similarity-space of architectures, with some creatures sharing some aspects of the global workspace with us to some degree, and others sharing others. In which case, many other animals turn out to share the publicly individuated property of phenomenal consciousness ("whatever property gives rise to the hard problem") to some significant degree; and this will, of course, be a matter of degree.

As we emphasized at the outset, however, phenomenal consciousness is *fundamentally* first-personal. It is the phenomenal concepts that one can apply introspectively in one's own case that define the subject-matter, and that are employed to generate the "hard" problems. And if, instead of formulating the question of animal consciousness using the public operationalized concept, we ask a question like, "Do other animals undergo states like *this* [feel of experience]?" then matters are not nearly so straightforward. Indeed, if the phenomenal concepts that delimit our subject-matter are first-person acquaintance-based indexicals, then what does it even *mean* to attribute phenomenal consciousness to another creature? This question can form our starting point for the remaining discussion.

Suppose one asks oneself, "Do other creatures, too, undergo *this* type of state?" where the indexical involved is an acquaintance-based phenomenal concept. What fixes the truth-condition for a correct answer? Plainly, as Kripke (1980)

emphasized, one is not using the indexical concept as a natural-kind concept. Although that concept does in fact refer to globally broadcast content (if the considerations adduced in sections 2 and 3 are correct), the concept is not employed with that intention. On the contrary, one means just to refer to the subjectively presented property that one is first-person acquainted with. One doesn't intend to think *through* the phenomenally conscious feelings in question to whatever they really are, to an underlying reality. And one plainly doesn't intend to refer to globally broadcast content as such. For one can deploy phenomenal concepts for one's own conscious states without even having the concepts of "global broadcasting" or "representational content." On the contrary, one intends to designate just one's own conscious feelings themselves.

Here is a related consideration: we can only know that our phenomenal concepts refer to globally broadcast content if we can *first* know that other people undergo phenomenally conscious states. Since the science of consciousness studies presupposes this, our first-person phenomenal concepts must somehow fix the truth-conditions involved in their projection to other people in advance of any discovery of what phenomenal consciousness really is.

Phenomenal concepts aren't natural-kind concepts, then. And since they are first-personal, neither are there any public norms guiding or constraining their use. (Of course we do *also* have public concepts for talking about phenomenal consciousness, including the concept phenomenal consciousness; but these are secondary to the first-person phenomenal concepts that fix the subject matter in question.)¹⁰ Moreover, since (as it turns out) there are no qualia, and neither are there any essentially first-personal properties of experience, the truth-condition of the thought that another person is undergoing a state of the same sort as *this* cannot be that the other person's mind, too, contains such properties. For there are none.

Although phenomenal concepts aren't used with the intention of referring to any natural kind, they might end up designating one nonetheless. Indeed, if we regard globally broadcast content as a natural kind, then phenomenal concepts *do* refer to a natural kind. But this is a kind that is human specific, since the functional properties of the global workspace include many components that are unique to human beings, including capacities for verbal report. Moreover, there is nothing in our use of these concepts that can make it the case that they designate some more-widespread but narrowly specified natural kind, such as attention, or working memory generally. Either our phenomenal concepts designate the whole enchilada—contents broadcast within the human-specific global workspace—or they fail to pick out any natural kind at all.

What, then, *does* fix the truth-conditions of thoughts attributing phenomenally conscious states to another creature? Consider this analogy. While visiting

10. See Balog (2009) for the distinction adverted to here between basic and nonbasic phenomenal concepts.

an area of a city to which one plans to relocate, one might exclaim, “*This* is the kind of neighborhood where we should live!” Although the type-indexical “*this*” here refers to a nonmental property, still its use is, in a sense, “private.” Plainly the intention is not to pick out a natural kind. Indeed, neighborhoods aren’t the sort of thing to *be* natural kinds (even though they may, in part, be composed of them). Nor are there any public norms guiding or constraining one’s use. Moreover, in the example as I imagine it, one might be incapable of articulating what it is about the neighborhood in question that feels so right.

Now suppose one asks whether another neighborhood in the city is also of *this* kind: what fixes the truth-conditions for a correct answer? I suggest the following: that if the dispositions-to-judge that underlie one’s initial thought were to be confronted with that other neighborhood in the same background circumstances of mood, weather, and so on, then they would issue in the same judgment. To project a private concept from one neighborhood to another requires holding the underlying dispositions-to-judge constant, and supposing them to be instantiated in that other neighborhood.

Similarly, then, with the phenomenal concept embedded in the question, “Do others undergo states of *this* sort?” Its projection into the mind of another person requires supposing that one’s current dispositions-to-judge are instantiated in the mind of the other. In fact, the truth condition for the thought, “Agent A undergoes states of *this* sort” (where *THIS* is a phenomenal concept) is as follows: “If the dispositions underlying my current use of the concept *THIS* were instantiated in the mind of agent A, then those dispositions would issue in a positive judgment that *this* is present.”

Where the agent in question is another adult human being, then attributions of states of *this* sort will generally be true (modulo factors like blindness and deafness), and one can know them to be true. This is not just because of the extensive similarities between our respective psychologies, knowable from a third-person perspective, or the underlying similarities in physical and neural organization. It is also because those other people, too, can express judgments much like my own. They can likewise conceive of zombies, and can readily become puzzled by the explanatory gap. This gives me every reason to think that they, too, can employ phenomenal concepts with similar overall properties to my own. So I can be confident that if the dispositions-to-judge underlying my phenomenal concept *THIS* were instantiated in those other minds, then those dispositions would issue in the positive verdict that a state of *this* sort is present.

In sum, then, to think that another creature is phenomenally conscious is to think that it undergoes states like *this* (where *THIS* expresses a first-person phenomenal concept). And that thought will be true just in case the dispositions-to-judge underlying my use of the concept *THIS* would issue in a positive verdict if they were to be instantiated in the mind of the creature in question.

When the creature one is thinking about has a mind significantly unlike one’s own, however, then attributions of phenomenally conscious states become

unevaluable, or truth-valueless. Suppose one wonders whether Macaque monkeys have phenomenally conscious experiences, for example. Macaque minds have much in common with our own. Specifically, they have attentional systems that control the entry of perceptual representations into a form of working memory, just as we do. But there are also important differences. (One might think of these as differences in the “global” component of the global workspace.) For their executive function capacities and planning capacities are much weaker than our own; they lack language altogether; it is unclear whether they engage in System 2 reflective thinking; and although there is some evidence that they can attribute some forms of mentality to other creatures, it is controversial whether or not they are capable of thinking about their own mental states at all (Carruthers 2019; Carruthers and Ritchie 2012; Carruthers and Williams 2019).

In supposing, then, that the dispositions-to-judge underlying one’s use of phenomenal concepts are instantiated in the mind of a Macaque, one might have to suppose that the mind of the Macaque is quite other than it actually is. For one has to suppose that it is capable of reflectively forming introspective indexical concepts about its own perceptual states. It might turn out that Macaques are capable of such a thing. But suppose I am right that they are not. In that case one cannot evaluate what would happen if Macaque minds *as they actually are* were to contain the dispositions-to-judge that underlie my own phenomenal concepts. If they did possess such dispositions, then their minds would be much more similar to the human mind than they actually are.

It turns out, then, that attempting to project first-personal phenomenal consciousness into the mind of another creature makes an important presupposition. It is presupposed that the mind of the creature in question could contain the dispositions-to-judge that underlie one’s use of phenomenal concepts without significantly altering what that mind is like. In connection with most nonhuman animals, that presupposition is false. And that means that the counterfactuals underlying the attribution of phenomenally conscious states to those animals can’t be evaluated. So there is no fact of the matter whether they have states *like this* or not. Attributions of phenomenal consciousness to them aren’t false, but are neither true nor false.

It seems that the question of phenomenal consciousness in other creatures turns, in part, on their capacity for higher-order thoughts. Does this mean that the theory of consciousness defended here is a version of higher-order-thought theory, after all? It does not. Phenomenal concepts are higher-order, of course. And it is phenomenal concepts that fix the domain of first-personal phenomenal consciousness. But what those concepts refer to is globally broadcast content. Indeed, phenomenal consciousness (at least in the human case) *is* globally broadcast content. That makes the theory a first-order one. But to ask about phenomenal consciousness in other creatures (to inquire about the extent of the domain picked out by our first-person phenomenal concepts) we have to project our phenomenal concepts into their minds. We have to ask whether, if the dispositions underlying

our phenomenal concepts were to be instantiated in those minds, our concepts would apply to any of their states. It is this that issues in indeterminacy.

It can be hard for people to accept that there is no fact of the matter whether or not most other animals have phenomenally conscious mental states. But I surmise that this is because they continue to assume that phenomenal consciousness is a *thing*, an essentially first-personal property of some sort that they are aware of in themselves. They assume that conscious states involve distinctive first-person properties that have a nature and/or existence over-and-above anything third-personally describable. But once one accepts that global broadcasting combined with the phenomenal-concept strategy fully explains everything there is to be explained, then there is no property remaining for us to inquire about its distribution across the animal kingdom. We can ask about the distribution of globally broadcast contents, of course. And the answer will be graded along multiple dimensions. But the only thing one *can* be asking if one insists on saying, “No, what about *phenomenal* consciousness? What about properties like *this* one?” is about how one’s dispositions to deploy phenomenal concepts would project into other creatures. That makes sense (and yields a determinate positive answer) for humans, but not for creatures with minds significantly unlike one’s own.

5. MARGINAL CASES

I have argued that the truth-conditions of thoughts ascribing phenomenal consciousness to other creatures involve counterfactuals that project the dispositions-to-judge underlying one’s own phenomenal concepts into the minds of those creatures. As a result, there are no facts of the matter about phenomenal consciousness in most species of animal. This is because the counterfactuals in question can’t be evaluated in minds that lack some of the capacities that underly and make possible those dispositions-to-judge.

There will, of course, be marginal and/or otherwise difficult cases, of which I will mention a few different sorts. In some instances, these will arise because we don’t know enough about the minds of the creatures in question. We don’t yet know, for example, the extent to which chimpanzees are capable of reflective thinking, nor whether they can entertain thoughts about their own current experiences. But a quite different source of uncertainty is that we might fail to know what capacities are really needed for a creature to possess phenomenal concepts like our own. Do such concepts require, not just capacities for higher-order thought and for reflective thinking, but also a capacity for language? Even if mental-state concepts aren’t generally language-dependent (as I believe), it could be that phenomenal concepts themselves are an exception. Perhaps introspectively applied type-indexical thoughts require a natural-language vehicle.

Other sorts of difficulty arise when one considers the emergence of phenomenal consciousness in human development from infancy to adulthood. Even if one

is quite clear about the cognitive capacities required for possession of phenomenal concepts, there remains the question of what counts as a *significant* change in the nature of the mind under consideration. Newborn infants are, of course, incapable of language, and they are likely incapable of reflective thinking or of entertaining higher-order thoughts about their own mental states. But their minds and brains contain nascent versions of all these capacities. So when one entertains the counterfactual, “If the dispositions-to-judge underlying my thought about *this* state were instantiated in the mind of an infant, then they would issue in the judgment that a state of *this* sort is present,” does this require us to suppose that the mind of the infant would be significantly unlike what it actually is, or not? If I imagine the infant’s mind to contain the capacities underlying those dispositions, am I imagining a very different kind of mind? It is hard to see what might settle the question either way.

Yet another sort of potential difficulty arises when we consider the possibility of phenomenal consciousness in artificial systems. To make the case as vivid as possible, imagine that we have reached the point where we can build androids that are cognitive isomorphs of adult humans. That is to say, they comprise systems that are functionally organized in the same manner as the human mind, and are capable of the same behaviors, including spontaneous expressions of puzzlement about the explanatory gap and the possibility of zombie versions of themselves. Would this be sufficient for the truth of the claim that the androids in question are phenomenally conscious? The answer may depend on one’s views about the metaphysics of dispositions (Choi and Fara 2018).

Recall that to ascribe phenomenal consciousness to another entity requires projecting one’s own phenomenal concepts into the mind of that entity. It requires us to evaluate a counterfactual conditional whose antecedent is, “If the dispositions-to-judge underlying my judgment about *this* state were instantiated in the mind of the other, then . . .” And then the question becomes: what has to be true for one of my cognitive dispositions to be instantiated in the mind of an android? If dispositions are identical with their categorical bases (Armstrong et al. 1996; Mackie 1977), then again the counterfactual becomes unevaluable at least, provided that the android’s cognitive system is realized in metal and silicon rather than living neurons. For on this view, my disposition-to-judge, having a categorical base involving neurons, cannot exist in the android’s mind, unless the android were to be differently physically constructed. But that requires us to suppose that the android is other than it actually is. So there is no evaluating whether the android, as it actually is, is phenomenally conscious.

Moreover, even if we retreat to a weaker view of the metaphysics of dispositions, merely identifying disposition *tokens*, rather than disposition *types*, with their categorical bases (Mumford 1998), the upshot is perhaps not so clear. For that requires us to decide whether one is projecting a token disposition-to-judge or a type disposition when one supposes that the dispositions underlying one’s use of phenomenal concepts are instantiated in the mind of the other agent.

Even if everyone were to agree that there is no fact of the matter concerning phenomenal consciousness in most species of animal, then, there would remain difficult and hard-to-answer questions at the margins. But are these important questions? Should we even attempt to answer them? I will now argue for a negative response.

6. WHO CARES ABOUT CONSCIOUSNESS?

Recall that on the view that I have been defending, there are no nonphysical qualia, and there are no essentially first-personal properties of experience. All that really exists is globally broadcast representational content. This can either be thought about as such, in the third person, or it can be thought about introspectively, employing acquaintance-based indexical concepts (that is, phenomenal concepts). So nothing magical entered the world with phenomenal consciousness. Nothing lit up. No spotlight got turned on. Broadcast contents just began entering into some novel functional relationships with other aspects of cognition.

This is why it doesn't matter that there is no fact of the matter about consciousness in most species of animal. What changes, as we look across species, are cognitive architectures that are more or less similar to our own along a number of cross-cutting dimensions (Carruthers 2019). Creatures can differ in their executive function abilities, in the forms of memory of which they are capable, in their planning abilities, in the richness of their conceptual repertoire, and much more. Many of these capacities can vary independently of one another across species. Generalist feeders like bears and raccoons, for example, will have extensive knowledge of the foodstuffs in their environments and their likely location and timing. But they may have quite limited planning abilities. Predators that mostly rely on a single source of prey, in contrast, like wolves and lions, might have quite limited conceptual knowledge of the environment, but significant planning capacities and strong impulse-inhibition abilities, enabling them to resist the urge to rush straight toward every potential source of food.

It seems that all mammals and birds, at any rate, have attentional mechanisms, as well as something resembling human working memory capacities (Carruthers 2015). Indeed, each of these things might be significantly more widespread even than this across the animal kingdom. What differs across species are the mechanisms that control the direction of attention, and the systems that comprise the global workspace, consuming and responding to the contents that enter working memory. As one surveys the resulting multidimensional similarity space, capacities get added, subtracted, enhanced, or reduced. But at no point does any special intrinsically first-personal property appear. At some stage (most likely in the primate or hominin lineage) there emerge a combination of abilities for higher-order thought, reflective thinking, and perhaps language, that together enable the formation of indexical acquaintance-based concepts about the contents of work-

ing memory. And at that point, the formation of phenomenal concepts becomes possible, and projecting them into the minds of one's conspecifics begins to make sense, resulting in true claims about the presence in those creatures of states like *this*. That is all there is to it.

Questions about the cognitive capacities of any given species of creature matter, of course, both for science, and for issues to do with humane treatment as well as animal husbandry (Dawkins 2017). We can (and should) inquire about the presence of negatively valenced sensory representations of bodily damage (pains) across species. And we can (and should) figure out what it is that animals prefer (either as a species or as individuals), for example by giving them options and seeing what they choose. But the fact that there is no fact of the matter about phenomenal consciousness ("states like *this*") in most species of animal doesn't matter for any of these things (Carruthers 2019). Moreover, there is no scientific mystery about the first emergence in nature of special first-person properties of experience, because there are none. And the fact that there is no fact of the matter isn't puzzling, either. It results merely from the fact that we cannot sensibly project our first-person phenomenal concepts into minds that are intrinsically incapable of forming or deploying them.

Given that facts about phenomenal consciousness don't matter, then questions about marginal or difficult cases of phenomenal consciousness don't matter, either. There might be some philosophical interest in figuring out whether dispositions are identical to their categorical bases. But this is not because it matters whether or not consciousness is in principle possible in a nonbiological android. Likewise, although there might be some philosophical interest in attempting to clarify the exact cognitive prerequisites for possession of phenomenal concepts, this isn't because it matters whether Macaques or chimpanzees are phenomenally conscious. The question of mentality in artificial systems is important, as are questions about the cognitive and affective capacities of Macaques and chimpanzees. But the question of phenomenal consciousness (states like *this*) in these entities isn't.

I have been arguing that the question whether (and which) other creatures are phenomenally conscious doesn't matter, on the grounds that no new first-personal property entered the world when states like *this* (phenomenal consciousness) did. Does that mean that I am committed to eliminativism or illusionism about consciousness (Frankish 2016)? It does not, or not straightforwardly. Since there are true claims about phenomenal consciousness, there are facts about phenomenal consciousness. Hence, in one good sense, phenomenal consciousness is real. When Mary emerges from her black-and-white room, for example, and thinks, "So *this* [content of experience] is what it is like to see red," she thinks something true. And likewise, I think something true when I think, while morbidly contemplating my own death, "When I die, all *this* [content of experience] will disappear." But the facts in question are *thin* rather than *thick*, tied to and individuated in terms of a particular mode of conceptualization, while having no reality beyond it.

7. CONCLUSION

It is time for us to stop thinking about phenomenal consciousness altogether. Certainly it shouldn't continue to attract the degree of scientific and philosophical attention it has received up to now. There are good reasons to believe that phenomenal consciousness (in humans) just *is* globally broadcast content, and that the so-called "hard" problem of consciousness arises merely from the distinctive set of concepts we employ when thinking about such contents in the first person. Third-person mentality matters, and understanding the multidimensional similarity-space of functionally organized minds across species matters. But first-personal phenomenal consciousness is just a distraction. We should stop caring about it.

ACKNOWLEDGMENTS

I am grateful to Jonathan Birch and two anonymous referees for their helpful comments on a previous version of this article, as well as to Tim Bayne and Nick Shea for related discussion.

REFERENCES

- Adam, K., E. Vogel, and E. Awh. 2017. "Clear Evidence for Item Limits in Visual Working Memory." *Cognitive Psychology* 97:79–97.
- Armstrong, D. M., C. Martin, and U. Place. 1996. *Dispositions: A Debate*. London: Routledge.
- Baars, B. 1988. *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.
- Ball, D. 2009. "There Are No Phenomenal Concepts." *Mind* 118:935–72.
- Balog, K. 1999. "Conceivability, Possibility, and the Mind-Body Problem." *Philosophical Review* 108:497–528.
- Balog, K. 2009. "Phenomenal Concepts." In *The Oxford Handbook of Philosophy of Mind*, edited by B. McLaughlin, A. Beckermann, and S. Walter. Oxford: Oxford University Press.
- Balog, K. 2012. "In Defense of the Phenomenal-Concept Strategy." *Philosophy and Phenomenological Research* 84:1–23.
- Blanke, O., T. Landis, and M. Seeck. 2000. "Electrical Cortical Stimulation of the Human Prefrontal Cortex Evokes Complex Visual Hallucinations." *Epilepsy and Behavior* 1:356–61.
- Block, N. 1995. "A Confusion about the Function of Consciousness." *Behavioral and Brain Sciences* 18:227–47.
- Block, N. 2007. "Consciousness, Accessibility, and the Mesh between Psychology and Neuroscience." *Behavioral and Brain Sciences* 30:481–99.
- Block, N. 2010. "Attention and Mental Paint." *Philosophical Issues* 20:23–63.
- Block, N. 2011. "Perceptual Consciousness Overflows Cognitive Access." *Trends in Cognitive Sciences* 12:567–75.
- Block, N. 2019. "What Is Wrong with the No-Report Paradigm and How to Fix It." *Trends in Cognitive Sciences* 23:1003–13.
- Carruthers, P. 2000. *Phenomenal Consciousness*. Cambridge: Cambridge University Press.
- Carruthers, P. 2015. *The Centered Mind: What the Science of Working Memory Shows Us about the Nature of Human Thought*. Oxford: Oxford University Press.

- Carruthers, P. 2017. "In Defense of First-Order Representationalism." *Journal of Consciousness Studies* 24:5–6, 74–87.
- Carruthers, P. 2018. "Valence and Value." *Philosophy and Phenomenological Research* 97:658–80.
- Carruthers, P. 2019. *Human and Animal Minds: The Consciousness Questions Laid to Rest*. Oxford: Oxford University Press.
- Carruthers, P., and J. B. Ritchie. 2012. "The Emergence of Metacognition: Affect and Uncertainty in Animals." In *Foundations of Metacognition*, edited by M. Beran, J. Brandl, J. Perner, and J. Proust. Oxford: Oxford University Press.
- Carruthers, P., and B. Veillet. 2007. "The Phenomenal-Concept Strategy." *Journal of Consciousness Studies* 14(9–10): 212–36.
- Carruthers, P., and B. Veillet. 2011. "The Case against Cognitive Phenomenology." In *Cognitive Phenomenology*, edited by T. Bayne and M. Montague. Oxford: Oxford University Press.
- Carruthers, P., and B. Veillet. 2017. "Consciousness Operationalized, a Debate Realigned." *Consciousness and Cognition* 55:79–90.
- Carruthers, P., and D. Williams. 2019. "Comparative Metacognition." *Animal Behavior and Cognition* 6:278–88.
- Chalmers, D. 1996. *The Conscious Mind*. Oxford: Oxford University Press.
- Choi, S., and M. Fara. 2018. "Dispositions." In *The Stanford Encyclopedia of Philosophy*, edited by E. Zalta. <https://plato.stanford.edu/archives/fall2018/entries/dispositions/>.
- Cohen, M. A., and D. Dennett. 2011. "Consciousness Cannot Be Separated from Function." *Trends in Cognitive Sciences* 15:358–64.
- Cohen, M. A., G. Alvarez, and K. Nakayama. 2011. "Natural-Scene Perception Requires Attention." *Psychological Science* 22:1165–72.
- Cohen, M. A., D. Dennett, and N. Kanwisher. 2016. "What Is the Bandwidth of Perceptual Experience?" *Trends in Cognitive Sciences* 20:324–35.
- Cutter, B., and M. Tye. 2011. "Tracking Representationalism and the Painfulness of Pain." *Philosophical Issues* 21:90–109.
- Dawkins, M. 2017. "Animal Welfare with and without Consciousness." *Journal of Zoology* 301:1–10.
- Dehaene, S. 2014. *Consciousness and the Brain*. New York: Viking Press.
- Ericsson, A., and W. Kintsch. 1995. "Long-Term Working Memory." *Psychological Review* 102:211–45.
- Frankish, K. 2016. "Illusionism as a Theory of Consciousness." *Journal of Consciousness Studies* 23(11–12): 11–39.
- Kapoor, V., A. Dwarakanath, S. Safavi, J. Werner, M. Besserve, T. Panagiotaropoulos, and N. Logothetis. 2020. "Decoding the Contents of Consciousness from Prefrontal Ensembles." *BioRxiv* preprint.
- Kripke, S. 1980. *Naming and Necessity*. Oxford, UK: Blackwell.
- Landman, R., H. Spekreijse, and V. Lamme. 2003. "Large Capacity Storage of Integrated Objects before Change Blindness." *Vision Research* 43:149–64.
- Lau, H., and D. Rosenthal. 2011. "Empirical Support for Higher-Order Theories of Conscious Awareness." *Trends in Cognitive Sciences* 15:365–73.
- Lewis-Peacock, J., A. Drysdale, K. Oberauer, and B. Postle. 2012. "Neural Evidence for a Distinction between Short-Term Memory and the Focus of Attention." *Journal of Cognitive Neuroscience* 24:61–79.
- Loar, B. 1990. "Phenomenal States." In *Philosophical Perspectives: Action Theory and Philosophy of Mind*, edited by J. Tomberlin. Atascadero, CA: Ridgeview Press.
- Ma, W., M. Husain, and P. Bays. 2014. "Changing Concepts of Working Memory." *Nature Neuroscience* 17:347–56.
- Mackie, J. L. 1977. "Dispositions, Grounds and Causes." *Synthese* 34:361–70.
- Mashour, G., P. Roelfsema, J-P. Changeux, and S. Dehaene. 2020. "Conscious Processing and the Global Neuronal Workspace Hypothesis." *Neuron* 105:776–98.
- Michel, M. 2017. "Methodological Artefacts in Consciousness Science." *Journal of Consciousness Studies* 24(11–12): 94–117.
- Mumford, S. 1998. *Dispositions*. Oxford: Oxford University Press.
- Panagiotaropoulos, T., G. Deco, V. Kapoor, and N. Logothetis. 2012. "Neuronal Discharges and Gamma Oscillations Explicitly Reflect Visual Consciousness in the Lateral Prefrontal Cortex." *Neuron* 74:924–35.
- Papineau, D. 2002. *Thinking about Consciousness*. Oxford: Oxford University Press.
- Prinz, J. 2012. *The Conscious Brain*. Oxford: Oxford University Press.

- Rosenthal, D. 2005. *Consciousness and Mind*. Oxford: Oxford University Press.
- Salti, M., S. Monto, L. Charles, J-R. King, L. Parkkonen, and S. Dehaene. 2015. "Distinct Cortical Codes and Temporal Dynamics for Conscious and Unconscious Percepts." *eLife*, 4, e05652.
- Sergent, C., V. Wyart, M. Babo-Rebelo, L. Cohen, L. Naccache, and C. Tallon-Baudry. 2013. "Cueing Attention after the Stimulus Is Gone Can Retrospectively Trigger Conscious Perception." *Current Biology* 23:150–55.
- Sligte, L., H. S. Scholte, and V. Lamme. 2008. "Are There Multiple Visual Short-Term Memory Stores?" *PLoS One*, 3, e1699.
- Sperling, G. 1960. "The Information Available in Brief Visual Presentations." *Psychological Monographs: General and Applied* 74:1–29.
- Thibault, L., R. van den Berg, P. Cavanagh, and C. Sergent. 2016. "Retrospective Attention Gates Discrete Conscious Access to Past Sensory Stimuli." *PLoS One*, 11, e0148504.
- Tononi, G. 2008. "Consciousness as Integrated Information: A Provisional Manifesto." *Biological Bulletin* 215, 216–42.
- Tononi, G., and C. Koch. 2015. "Consciousness: Here, There, and Everywhere?" *Philosophical Transactions of the Royal Society B*, 370:20140167.
- Trübutschek, D., S. Marti, A. Ojeda, J-R. King, Y. Mi, M. Tsodyks, and S. Dehaene. 2018. "A Theory of Working Memory without Consciousness or Sustained Activity." *eLife*, 6, e23871.
- Tsubomi, H., K. Fukuda, K. Watanabe, and E. Vogel. 2013. "Neural Limits to Representing Objects Still within View." *Journal of Neuroscience* 33:8257–63.
- Tye, M. 1995. *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
- Tye, M. 2000. *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.
- Tye, M. 2009. *Consciousness Revisited: Materialism without Phenomenal Concepts*. Cambridge, MA: MIT Press.
- Veillet, B. 2012. "In Defense of Phenomenal Concepts." *Philosophical Papers* 41:97–127.
- Vignal, J., P. Chauvel, and E. Halgren. 2000. "Localized Face Processing by the Human Prefrontal Cortex: Stimulation-Evoked Hallucinations for Faces." *Cognitive Neuropsychology* 17:281–91.
- Wu, C-C., and J. Wolfe. 2018. "A New Multiple Object Awareness Paradigm Shows That Imperfect Knowledge of Object Location Is Still Knowledge." *Current Biology* 28:3430–34.
- Xia, Y., Y. Morimoto, and Y. Noguchi. 2016. "Retrospective Triggering of Conscious Perception by an Interstimulus Interaction." *Journal of Vision* 16:3.