



Consciousness operationalized, a debate realigned



Peter Carruthers^{a,*}, Bénédicte Veillet^b

^a Department of Philosophy, University of Maryland, United States

^b Department of Philosophy, University of Michigan–Flint, United States

ARTICLE INFO

Keywords:

Access consciousness
Amodal experience
Cognitive phenomenology
Hard problem
Phenomenal consciousness
Zombie

ABSTRACT

This paper revisits the debate about cognitive phenomenology. It elaborates, defends, and improves on our earlier proposal for resolving that debate, according to which the test for irreducible phenomenology is the presence of explanatory gaps. After showing how proposals like ours have been misunderstood or misused by others, we deploy our operationalization to argue that the correct way to align the debate over cognitive phenomenology is not between sensory and (alleged) cognitive phenomenology, but rather between non-conceptual and (alleged) conceptual or propositional phenomenology. In doing so we defend three varieties of non-sensory (amodal)¹ non-conceptual phenomenology: valence, a sense of approximate number, and a sense of elapsed time.

1. Introduction

The debate over cognitive phenomenology is often construed as a debate about whether thoughts and concepts have phenomenal properties that are irreducible to broadly sensory ones. (There may be a number of ways of characterizing what counts as “broadly sensory”. We return to this question in Section 3.5.) On the one side, so-called *conservatives* about cognitive phenomenology claim that only sensory states are phenomenally conscious (Carruthers, 2000; Jackendoff, 1987, 2012; Prinz, 2012; Tye, 1995, 2000). By contrast, so-called *liberals* maintain that the phenomenology associated with thoughts and concepts can’t be fully reduced to broadly sensory phenomenology (Strawson, 1994, 2011; Siewert, 1998, 2011; Pitt, 2004; Chudnoff, 2015a,b; Kriegel, 2015; McClelland, 2016).

Liberal arguments take a number of forms, but many of them appeal to intuition or to the existence of phenomenal contrasts. For example, they invite us to reflect on the difference in experience between a French speaker and a non-French speaker listening to one and the same sentence (that is, with and without understanding). Conservatives, however, typically agree that there are such differences, while pointing out that there will be many disparities in *sensory* phenomenology caused by underlying differences in understanding. For example, linguistic competence will have an impact on how the phonology of the sentence is represented, and understanding may be accompanied by a suite of imagistic and affective differences, all of which can count as sensory. The phenomenal differences between the two speakers listening to the same sentence, the conservative concludes, are ultimately reducible to broadly sensory phenomenology.

* Corresponding author.

E-mail address: pcarruth@umd.edu (P. Carruthers).

¹ Note that the term “amodal” admits of two quite different uses in cognitive science. One (the one we intend) refers to processes or representations that are neither sensory-specific nor multi-sensory (or “multi-modal”) in nature. We assume throughout that concepts are amodal in this sense. The other use is the one that figures in the phrase “amodal completion” in vision science. This refers to the construction by low-level visual processes of an imaginary boundary of a partially-occluded object. The contrasting sort of completion – “modal completion” – refers to the construction of an imaginary boundary of the presumed occluding object. Both are exemplified in the famous Kanizsa triangles. Moreover, both are modality-specific (specific to the visual system) in our intended sense of “modal”.

The mere existence of phenomenal contrasts between understanding and not understanding a sentence doesn't settle the question of cognitive phenomenology, then. Neither do straightforward appeals to intuition, which aim to support a liberal position by urging us to *introspect* on the phenomenology associated with thoughts or concepts (such as the phenomenology associated with the thought that I have found my keys). Instead, the arguments on either side are best construed as forms of inference to the best explanation (Bayne & Montague, 2011, 22). What must be shown is not simply that there *are* contrasts of the sort described, or that our thoughts are *associated with* phenomenology, but rather that liberals can provide a better explanation of these facts than can conservatives. Indeed, as Chudnoff (2015a) points out, the liberal's claim about the irreducibility of cognitive phenomenology is actually a "(i) logically complex, (ii) generalization about (iii) possible (iv) explanatory relations" (24). It seems quite unlikely that the truth of such a generalization could be revealed by the mere existence of phenomenal contrasts or by introspection alone.

Given that what is at stake is the best explanation of the contrasting cases, or of the deliverances of introspection, it is worth stressing that this makes the conservative position the default. Since *everyone* accepts the existence of sensory phenomenology, it is simpler to maintain that there is *only* sensory phenomenology, in the absence of something that cannot be explained adequately in those terms. The point here is an entirely general one about inference to the best explanation. Explanations, in general, are better if they are simpler, requiring us to postulate fewer types of entity or kinds of property. Since conservatives and liberals offer competing explanations of the phenomenology distinctive of thought, and since everyone accepts the sorts of sensory phenomenology that conservatives appeal to in their accounts, the onus is on liberals to make their case. To be a liberal means accepting a novel set of phenomenal properties, and this needs justification. Liberals *claim* to have found something beyond the sensory phenomenology that everyone accepts, of course, but they have failed to convince their opponents of what that might be.

In Carruthers and Veillet (2011) we proposed a way forward, inspired in part by Bayne (2009). We suggested that an independent test for the presence of phenomenology of any specific sort is to consider whether it gives rise to (or at least *seems* to give rise to) thought experiments of the "hard problem" sort (Chalmers, 1996), thus setting up at least the appearance of an "explanatory gap" between physical and functional facts, on the one hand, and the properties in question, on the other. We argued (and continue to maintain) that our proposed operationalization is a reasonable one. For what could be the point of describing some state or property as *phenomenally* conscious if it didn't give rise to any hard-problem-type thought experiments? Indeed, it seems us that it is the mark of phenomenal consciousness (as opposed to mere access-consciousness) that it naturally raises (or at least seems to raise) a set of *hard epistemic problems* – problems that other aspects of our mental lives don't give rise to (including any aspects that are merely access-conscious but not phenomenally conscious). This doesn't require us to claim that the explanatory gap is real or insuperable, of course; let alone that it supports any form of metaphysical dualism. In fact, our proposal was intended as a test for the presence of phenomenal consciousness that could be neutral between all of the opposed positions about the nature of consciousness.

At around the same time a number of philosophers arrived at very similar ideas (Horgan, 2011; Kriegel, 2009), and in the ensuing years other contributors to the debate have embraced and further developed this kind of proposal, suggesting that the presence or absence hard-problem-type thought experiments can be used as a diagnostic tool for the presence or absence of any disputed sort of phenomenology (Kriegel, 2015; McClelland, 2016). We think, however (and will argue here), that in some cases the test has been misapplied, and that when used properly it leads to a more nuanced position – indeed, to one that vindicates neither conservatism nor liberalism as these positions have traditionally been understood in the existing debate.

In Section 2, we revisit our proposal. We begin by responding to recent criticisms aimed directly at our earlier use of hard-problem thought experiments (McClelland, 2016). We then discuss the further use made of such thought-experiments by Horgan (2011) and McClelland (2016), who rely on them to defend liberalism about cognitive phenomenology (specifically, the phenomenology of understanding what someone is saying). Our critique will find a new use for the familiar distinction between *access-consciousness* and *phenomenal consciousness* (Block, 1995).

In Section 3 we use our operationalization of phenomenal consciousness in constructive mode, arguing for a realignment in the debate over cognitive phenomenology. Conservatives are normally said to endorse only *sensory* phenomenology (understood broadly, to include proprioceptive and interoceptive sensations, as well as affective states in general), whereas liberals are said to believe that thoughts and other states with conceptual or propositional content have phenomenology that is irreducible to anything broadly sensory. We argue, in contrast, that there are at least three sorts of amodal, non-sensory, but nevertheless non-conceptual kinds of phenomenal state, including (a) the valence component common to all forms of affective state, (b) our experience of approximate number (generally called "numerosity"), and (c) our experiences of time and temporal distance. We suggest that all and only non-conceptual states are phenomenally conscious (when access-conscious), no matter whether those states are sensory or not.

In Section 4 we turn to what is arguably the most sophisticated and challenging argument for a strongly liberal position, due to Kriegel (2015). He imagines the case of Zoe, who not only (like a zombie) lacks sensory phenomenology, but lacks sensory states altogether, in all sensory modalities. However, Zoe is still said to be capable of cognitive phenomenology resulting from her purely intellectual work as an accomplished mathematician. We argue that the forms of phenomenology that might plausibly remain in the case of Zoe can successfully be explained in terms of the kinds of amodal non-conceptual states defended in Section 3. Section 5 is our concluding discussion.

Before we get to our main discussion, however, we should say something further about the conceptual/non-conceptual distinction, since this will figure prominently in what follows. We take the basic contrast to be one between representations that involve categorical boundaries of some sort (or are "chunked"), and those that are fine-grained or continuous (or "analog") in nature. This

way of drawing the distinction between conceptual and non-conceptual content is pretty standard in the literature (Bermúdez, 2015; Tye, 2000), and has been familiar since at least Peacocke (1992).²

Non-conceptual content can also be characterized contrastively, however, as representational content that is *not conceptual* in nature. And much then turns on how concepts themselves are understood. Many philosophers follow Evans (1982) in placing strong constraints on genuine concept possession. Specifically, many think that concepts are subject to the “Generality Constraint”. Roughly stated, this requires that if a creature is to possess a given concept, then it must be capable of thinking thoughts that combine that concept with *any other* concept (of appropriate logical form) that the creature also possesses. So a creature that possesses a singular concept of Juliette and a singular concept for the sun (together with a concept of identity) must be capable of entertaining the thought, *Juliette is the sun*, for example. As should be obvious, this is very likely a constraint that the representational states of non-human animals fail to meet. Animals are thus, at best, said to be capable of employing “proto-concepts” and entertaining “proto-thoughts” (Bermúdez, 2003; Dummett, 1994).

One of us has argued that the Generality Constraint is much too strong, however (Carruthers, 2009), and that the class of non-conceptual states that it carves out is much too broad (including the belief-like states of non-human animals). It needs to be accepted, of course, that concepts are the building blocks of thoughts, since thoughts are compositionally structured out of concepts. But one can accept this without requiring that those building-blocks should be combinable by the thinker in *any* logically permissible combination. This would be to require complete flexibility of thought as a condition for entertaining thoughts at all. Instead, one can require only that concept possession calls for one to be capable of combining that concept together with *some* other concepts that one has. In what follows, then, we will construe conceptual content more broadly than a strong Generality Constraint would permit, allowing that animals and infants are capable of thought. And we thus understand non-conceptual content more narrowly, as the sort of content that is more fine-grained than can figure in the content of a creature’s thoughts.

2. Partial zombies

In Carruthers and Veillet (2011) we used our proposal to argue for a conservative position about cognitive phenomenology. We set out to show that the phenomenology associated with thoughts and conceptual states doesn’t give rise to the sorts of hard-problem thought experiments that indicate the presence of irreducible phenomenology. In that paper we mostly relied on content-inversion thought experiments. For instance, we pointed out that someone who is looking at a red tomato can think, “*This* experience [the reddish one] might not have represented anything red, and could have been reliably produced by green things instead.” In contrast, we argued, someone who is viewing a duck-rabbit figure as a duck (deploying in her experience the concept *duck*) cannot coherently think, “*This* experience [the seeing-as-a-duck one] might not have represented duckhood, and could have been reliably caused by something else instead.”

McClelland (2016) takes issue with our use of content-inversion thought experiments, however. He suggests that by arguing as we did, we might have biased the case against cognitive phenomenology. Following Bayne (2009), McClelland notes that it is far from clear that all forms of experience that are indisputably phenomenally conscious admit of content-inversion thought experiments. Experiences of color and sound do. But it is by no means obvious that the same holds for visual experiences of primary qualities like shape. McClelland argues, in consequence, that the test for phenomenal consciousness should be the availability of zombie or partial-zombie thought experiments. For these apply to experiences of primary qualities like shape just as much as they apply to experiences of secondary qualities like color.

We think that McClelland (2016) is correct on this point. While the conceivability (or apparent conceivability) of content-inversion thought experiments may be a reliable diagnostic tool for the *presence* of phenomenal consciousness, *failures* of such apparent-conceivability aren’t a reliable sign of the *absence* of phenomenal consciousness. (The same is true of what-Mary-doesn’t-know thought experiments, which likewise only obviously apply to experiences of secondary qualities.) Of course (as McClelland himself grants) this doesn’t undermine the spirit of our proposal, which is that irreducible phenomenology should give rise to hard-problem thought experiments. For some of these thought experiments – more specifically zombie thought-experiments – have more general application. Indeed, they arguably coincide with any form of experience that gives rise to at least the appearance of an explanatory gap. And as we noted above, what would be the point of describing some mental property as “phenomenally conscious” if it didn’t even seem to create such a gap? Accordingly, in what follows, we too will make use of zombie and partial-zombie thought experiments. Our question in what remains of this section is whether those who have used these thought-experiments – Horgan (2011) and McClelland himself (2016) – are right that they support a liberal position.

Let us first consider Horgan’s (2011) version of the claim that the presence of irreducible phenomenology should lead to epistemic puzzles and gaps. Here is how he puts it, “*the robust conceivability of certain kinds of zombie scenarios can serve as a criterion for the existence of certain kinds of phenomenal character*” (61, emphasis in the original). More specifically, he argues that to show that a certain sort of cognitive phenomenology is irreducible to other (broadly sensory) phenomenology, we should ask whether we can robustly conceive of a certain kind of partial zombie. We are happy to agree. The right kind of zombie is specifically one who: (1) is a complete functional duplicate of an ordinary person, (2) has the same *sensory* phenomenology as an ordinary person, but who (3) is nonetheless missing the phenomenology typically associated with belief or thought (Horgan, 2011, 61). The idea is this: if the conservative is right and cognitive phenomenology is indeed reducible to broadly sensory phenomenology, then it should be

² Those who deny the existence of non-conceptual content (McDowell, 1994) could recast what we say in terms of course-grained versus fine-grained indexical judgments.

impossible for someone to have all our ordinary non-cognitive (sensory) phenomenology and yet be missing cognitive phenomenology. In other words, the kind of zombie Horgan urges us to imagine would be impossible; and we would thus be unable to robustly conceive of such a zombie.

To illustrate, consider Horgan's Andy₃, who is stipulated to suffer from "language-understanding absence partial zombie disorder" (2011, 69), all the while being a duplicate of regular, ordinary, Andy₁ when it comes to both functional and sensory/affective phenomenal properties. As such, Andy₃ follows (and would follow) "an exactly similar lifelong behavioral trajectory through the world" to the one that Andy₁ does (70). However, from a first-person perspective things are rather different for Andy₃: "the sounds and marks that Andy₁ experiences as intelligible language are always experienced by Andy₃ as meaningless noises and squiggles – even when he produces them himself" (70). Of course, since Andy₃ is otherwise a duplicate of Andy₁, he will still manage to communicate, and it will appear to others as though he is uttering meaningful sentences: it is just that *he* will experience all of those sentences (whether uttered by himself or others) as *meaningless*.

There are a number of problems with Horgan's thought experiment (which we take to be representative of the further examples he develops). Indeed, there is ample reason to doubt that Andy₃ could differ in phenomenology from Andy₁, given that they are causal/functional duplicates. Since it is stipulated that Andy₁ genuinely understands English, he will certainly experience English sentences as meaningful. And since they are causal/functional duplicates, there will be no differences between Andy₁ and Andy₃ in terms of phonological, syntactic, or semantic processing; both will be able to identify word boundaries in ways that are equally rich and determinate; and both will have the same emotional and imaginative responses to the strings of sounds they hear. Moreover, both will immediately draw inferences from what is said, will often know immediately what to say in reply, will be disposed to report what has been said if asked to do so, and so on. In addition, each will be disposed to deny (even to a trusted therapist) that there is anything untoward about the way speech feels to him.

As we continue to detail what is required (in terms of causal/functional structure) for Andy₁ to understand English, it becomes increasingly hard to maintain that we can *robustly conceive* of all that causal/functional processing being identical in each of them *except* that Andy₃ experiences English sentences as *meaningless*. After all, how can Andy₃ experience a string of words as meaningless, while nevertheless coming to believe (by the same causal route as Andy₁) that he knows exactly what was just said to him, and while being able to report it in other words?

The case here turns out to be rather similar to our duck-rabbit example (discussed earlier in this section and more fully in Carruthers & Veillet, 2011). It seems impossible that two individuals could be causally/functionally identical (in terms of both visual processing and concept deployment) while one experiences the ambiguous figure as *a duck* and the other experiences it as *a rabbit*. In the same way, it seems impossible that Andy₁ and Andy₃ could truly be causally/functionally identical (in terms of linguistic processing and concept deployment) while Andy₁ experiences a spoken English sentence as *meaningful* and Andy₃ experiences it as *meaningless*.

Horgan's (2011) mistake, in our view, is that he fails to appreciate the relevance to his thought experiment of the distinction between access-consciousness and phenomenal consciousness.³ That there is at least a conceptual distinction between the two isn't in dispute.⁴ To say that a mental state is access-conscious is to say (in modern parlance) that it is *globally broadcast* to a wide range of cognitive systems – for forming memories, for reporting in speech, for action-planning, for use in executive decision making, and to systems that create full-blown emotional responses, among others (Dehaene, 2014). Phenomenal consciousness, in contrast, is best characterized in terms of its aptness to give rise to hard-problem thought experiments (or so we suggest). Now notice that access-consciousness is implicit in the very idea of a zombie: a zombie is supposed to be a creature that is physically and functionally indistinguishable from a normal person. That means that it, too, must have content-bearing states that are globally available to just the same range of functions and cognitive processes as a normal person.

Horgan, however, doesn't seem to fully recognize that a partial zombie's state of understanding – the state of grasping the meaning of a sentence – would have to be access-conscious in order for it to qualify as a partial *zombie*. After all, the meaning of the sentence is access-conscious for regular, ordinary, Andy₁. So the meaning of the heard sentence would have to be globally broadcast to Andy₃'s systems for memory, for planning, for verbal report, and so on, thereby having a direct impact on those systems of just the sort that happens in a Andy₁. In consequence, Andy₃ would be *aware* of the meaning in a purely functional, access-conscious, sense of "aware". Given the identity in sensory experience between Andy₃ and his normal counterpart Andy₁, and given that both are immediately aware (in the access-conscious sense) of the same meaning, what else could Andy₃ possibly lack?⁵

A similar failing seems to us to infect McClelland's (2016) defense of liberalism, which is closely modelled on Horgan's. He, too, imagines a partial-zombie scenario, but he focuses on the change that takes place when either the partial-zombie or its twin come to

³ It isn't just liberals who tend to lose sight of this distinction. The same thing can happen to conservatives. For example, Jackendoff (2012) thinks that his defense of a conservative position commits him to claiming that "meanings are ... in a part of our minds that we can't access" (81); or more colorfully: "The meaning itself stays backstage" (84). This is plainly a mistake. The meanings of the sentences we understand are obviously access-conscious (they can be recalled, can be reported, can enter into planning, can give rise to emotional reactions, and so on), even if they aren't *phenomenally* conscious.

⁴ What is disputed is whether phenomenal-consciousness can "overflow" the contents of access-consciousness in Sperling-style experiments (Sperling, 1960), in which people report only four items from a briefly presented complex array while claiming to have seen the entire display, but can (roughly speaking) report any four items if cued to their position following stimulus off-set. Block (1995, 2007, 2011, 2014) argues that phenomenal consciousness overflows access-consciousness in these sorts of cases, whereas a variety of critics argue that it doesn't (Cohen & Dennett, 2011; Kouider, de Gardelle, Sackur, & Dupoux, 2010; Phillips, 2011; Stazicker, 2011). We will set aside this debate here, since it isn't relevant to our topic. For the views of one of the present authors, see Carruthers (2015a).

⁵ In fact, there may well be a small residue of phenomenally conscious properties that we can imagine to be absent in the case of Andy₃ without compromising the identity between his sensory states and those of Andy₁, having to do with the non-conceptual (but amodal) content involved in *fluent cognitive processing* of a stimulus. This point will emerge more clearly in Section 4 when we discuss Krieger's (2015) example of the case of Zoe.

understand a proposition meaning *ducks are cute* for the first time. In such circumstances, McClelland writes, “you come to experience something that your twin does not: you come to experience the proposition *that ducks are cute*” (548). But McClelland’s use of the term “experience”, here, begs the question, given that phenomenal consciousness is often characterized *as* experience (Block, 1995). The most that McClelland is entitled to assume without further argument is that the proposition *ducks are cute* becomes access-conscious when you understand “ducks are cute”. Everyone can agree you come to be *aware* of that proposition (in the access-conscious sense) when you understand the sentence. But your zombie twin, too, would come to be directly aware of the proposition that ducks are cute, given that the content in question is globally broadcast and received as input by the usual suite of receiving systems (for remembering, reporting, drawing inferences, and so on). And it is far from obvious that there is anything remaining for your zombie twin to lack. Indeed, we suggest that one cannot successfully think anything of the form, “There could be someone like me in all sensory respects, and in whom the proposition that ducks are cute is immediately available in the usual way for remembering, reporting, drawing inferences, and so on, but in whom *this* feeling [attaching to the proposition, *ducks are cute*] is absent.” For which feeling is that?

McClelland (2016) might reply that we have mischaracterized the intent behind his partial-zombie scenario. Specifically, although physically, functionally, and sensory-state identical, his zombie-twin might be supposed to *lack* direct access to the proposition *ducks are cute*. It is the content of that proposition itself that is missing in the case of the partial-zombie, and that makes his overall experience so different, it might be said. But this is to undercut the very idea of a zombie. (Recall that a zombie, by stipulation, is physically and functionally identical to us; that is, identical in respect of *access-consciousness*.) For there is no theory of content – whether causal-informational (Dretske, 1988; Fodor, 1990), or teleosemantic (Millikan, 1984; Papineau, 1987), or inferential-role (Block, 1986), or some combination thereof (Botterill & Carruthers, 1999; McGinn, 1982) – according to which all the physical and functional facts can be held fixed (including causal and counterfactual facts) and yet content can be different or absent. And notice, too, that content-facts *can* be (and are) held fixed in standard zombie thought experiments about sensory qualities. One can seemingly think, “There could be a creature physically just like me, who has access-conscious states representing red [i.e. that have the content *red*], but in whom *this* feel of my experience of the tomato is absent.”

One last point. We have sometimes encountered objectors who respond to our argument as follows: “How could one *experience* a sentence as meaning that ducks are cute unless there is cognitive phenomenology? For after all, experiences are phenomenal states, aren’t they?” But this is to miss the point that access-conscious states might have both phenomenal and non-phenomenal components. We grant that when one understands a sentence one is aware, both of its sensory properties (pitch, loudness, phonology, and so forth) and its semantic ones. That is, in such cases one-and-the-same globally broadcast event-file will comprise both sensory/non-conceptual representations and propositional/conceptual ones (Carruthers, 2015b). While the former contribute to the phenomenal properties of one’s conscious state, the latter might not. So, yes, the experience of a sentence as meaning that ducks are cute is a phenomenal state; but it doesn’t follow that the meaning-component of that experience has phenomenal properties, nor that it makes a constitutive, irreducible, contribution to the phenomenal properties of the entire hearing-event.

Our interim conclusion is that while Horgan (2011) and McClelland (2016) have improved on Carruthers and Veillet (2011) by using zombie thought experiments rather than content-inversion ones to get at the question of an explanatory gap, they have failed to demonstrate that there is any such gap between physical and functional facts, on the one hand, and cognitive states like those involved in understanding a sentence, on the other. In consequence, the simplest hypothesis wins: a conservative position about cognitive phenomenology should be our preferred option.

3. Beyond sensory phenomenology

In the present section we explore the consequences of our proposal for aspects of phenomenally conscious experience that are rarely mentioned in the context of the current debate. We argue that our operationalization shows that contents distinctive of affective states (including the valence component of such states) and our sense of approximate number, as well as those integral to our sense of time, are phenomenally conscious. This suggests, we believe, that all non-conceptual states are phenomenally conscious (at least, when they are “globally broadcast” and hence access-conscious). Note that some, at least, of these non-conceptual contents are non-sensory (amodal) ones. This is important, because there has been a persistent tendency in the literature to equate the denial of cognitive phenomenology with the claim that only sensory and sensory-like (e.g. imagistic or emotional) states are phenomenally conscious (Bayne & Montague, 2011; Kriegel, 2015). This is false, we think. The domain of the phenomenal is wider than the sensory (albeit not including the conceptual). We will come back to this point at the end of the section.

3.1. Emotion

Most participants in this debate have allowed that emotions are phenomenally conscious, or have phenomenally conscious components, at least. Although the essence of emotion is much disputed (Ellsworth & Scherer, 2003; Prinz, 2004; Russell, 2003), all will allow that emotions are generally associated with appraisals of environmental stimuli (as dangerous, for example), action-tendencies (running away, in the case of fear), facial expressions, bodily arousal (increased heart rate, breathing rate, and so on), and valence (often described in terms of pleasure or displeasure). Facial expressions and bodily-arousal properties will give rise to non-conceptual proprioceptive and interoceptive sensory experiences which (when attended to and globally broadcast) will thereby become phenomenally conscious. Conservatives, however, will deny that seeing something *as dangerous* (a conceptual state) is itself phenomenally conscious, although of course it can *cause* states that are phenomenally conscious (experiences of increased heart rate, sweaty palms, and so on).

While no new issues for debates about phenomenal consciousness are raised by the sensory and conceptual aspects of emotion, valence is another matter. For this, as we will see, is neither sensory nor conceptual.

3.2. Valence

Valence is an aspect of all emotional states. Fear is negatively valenced, for example, whereas joy is positively valenced. But it is also a component in affective states more generally, including pain, orgasm, sensuous touch, the experience of drinking when thirsty or eating when hungry, the experience of conversing with a friend, and more. Increasingly, valence is seen by cognitive scientists as the “common currency” underlying human and animal decision-making (Levy & Glimcher, 2012; Ruff & Fehr, 2014), enabling us to trade off goods and bads against one another even when (on a conceptual level) they seem incommensurable.⁶

Note that if the common-currency view of valence is correct, then valence cannot be a sensory or sensory-specific state. For the negative valence of pain will be the same as the negative valence of grief, which is the same as the negative valence contained in fear or disgust. Indeed, not only can grief and other forms of social suffering be blunted by taking Tylenol, just as can physical pain (Lieberman, 2013; Lieberman & Eisenberger, 2009), but so, too, is pleasure blunted by the same drugs (Durso, Luttrell, & Way, 2015). Moreover, valence is not itself a conceptual state, although it can be caused by cognitive/conceptual appraisals of a thing or situation. It is continuously graded (in this respect like sensory states), providing a fine-grained representation of value that grounds one’s conceptual judgments of *good* and *bad* in something like the way that fine-grained representations of color ground one’s judgments of color (Carruthers, 2017). It seems that valence is not only amodal (or non-sensory), but also non-conceptual. It provides a first test-case, then, of the widespread assumption that denying cognitive (conceptual) phenomenology commits one to the view that only sensory states are phenomenally conscious. It also provides a test for our thesis that all non-conceptual contents (when access-conscious) are phenomenally conscious.

Consider the negative-valence component of pain: the *hurtfulness* of pain. As is now well known, pain comprises both sensory and valuational aspects, underlain by partially distinct neural pathways. The sensory component can remain the same while its unpleasantness either increases or decreases. Indeed, some forms of morphine can leave the sensory aspect of pain unchanged, while removing the negative valence altogether. People say that their pain sensations remain just as they were, but they no longer *feel bad*. The same can happen when people undergo anterior cingulotomy as a treatment for chronic pain (Wilkinson, Davidson, & Davidson, 1999).

The sensory component of pain is phenomenally conscious, of course. But so too, it seems, is the negative-valence component. One’s experience is phenomenally very different when a pain sensation is no longer hurtful (even if the sensation itself is unchanged). However, a familiar problem rears its head at this point: how can we tell that the contribution made by valence to the phenomenal qualities of pain isn’t reducible to sensory properties? How can we tell that the phenomenology associated with valence is irreducible to other phenomenology? For if one is no longer bothered by one’s pain sensations one will attend to them less, one will no longer grimace and tense one’s muscles whenever the sensation returns, and so on – with consequent effects on one’s overall phenomenal experience.

At this point our proposed operationalization of phenomenal consciousness can help. One can surely imagine someone just like oneself, who experiences sensations of pain and who very much wants (conceptually) that those sensations should cease, but whose pains don’t *hurt*.⁷ Such a person would, indeed, grimace and tense whenever a sensation of pain returns, since she doesn’t want it to occur. But her phenomenal experience would surely be very different from ours. In fact, it seems one can coherently think the thought, “There could be someone just like me in all respects, who feels pain sensations and who behaves just as I do when she has them because she wants them to cease, but in whom *this* component of the experience of pain [its hurtfulness] is missing.”

Indeed, it seems that the full gamut of epistemic puzzles raised by phenomenal consciousness arise also in the case of hurtfulness (and valence in general).⁸ For instance, one can run the equivalent of Jackson’s color-deprived-Mary thought experiment (Jackson, 1982, 1986). Imagine someone with congenital pain asymbolia (who feels pain sensations but has never been bothered by them) who becomes a famous psychologist and neuroscientist. She learns everything there is to know about pain (its causes, functions, and physical realization). Surely, one might think, she would nevertheless learn something *new* if she were cured of her asymbolia and became capable of being hurt by her pains for the first time. (Compare how color-deprived Mary is supposed to learn something new about color vision when leaving her black-and-white room for the first time.) And note that what she learns isn’t what it is like to grimace and tense her muscles, since we can suppose she has already experienced that for other reasons. When she jabs herself with a pin for the first time following her cure she might exclaim, “So *this* is what the hurtfulness of pain is like!”

If phenomenal consciousness is operationalized as whatever gives rise to “hard problem” thought experiments, then there is a strong case for saying that valence is (constitutively/irreducibly) phenomenally conscious. And since valence is both non-conceptual

⁶ Valence-processing across all domains and types of affect appears to be underlain by a single (albeit multicomponent) neurobiological network, involving not just subcortical evaluative regions in the basal ganglia, but also the anterior insula and anterior cingulate, together especially with orbitofrontal and ventromedial prefrontal cortex (Bartra, McGuire, & Kable, 2013; FitzGerald, Seymour, & Dolan, 2009; Leknes & Tracey, 2008; Plassmann, O’Doherty, & Rangel, 2010). The latter regions are the primary projection areas for valence signals in the cortex, and seem to be the ones implicated in prospective reasoning and decision making. When these regions are damaged, people can still make perfectly sensible *judgments* about what is good, or about what should be done, but their actual decision-making goes all to pieces (Bechara, Damasio, & Damasio, 2000; Damasio, 1994).

⁷ We here assume the falsity of desire-theories of the hurtfulness of pain (Armstrong, 1962; Tye, 1995). For a powerful critique, see Bain (2014).

⁸ Recall from Section 2 that although only zombie-style thought experiments are reliable indicators of the *absence* of phenomenal consciousness (as McClelland, 2016, argues), the remaining hard-problem-type thought experiments can be reliable indicators of the *presence* of phenomenal consciousness.

and non-sensory (amodal) in nature, it follows that some phenomenally conscious states are neither sensory nor conceptual. We have also taken a first step toward our target generalization: that *all* non-conceptual states (when access-conscious) are phenomenally conscious.

3.3. Numerosity

As is now well-known, humans and many other animals can represent (and add, divide, and subtract) the numerosity (approximate number) of large sets (Barth, Kanwisher, & Spelke, 2003; Gallistel, Mark, King, & Lantham, 2001; Izard, Sann, Spelke, & Streri, 2009; Jordan, MacClean, & Brannon, 2008). For instance, if someone taps out a rhythm as fast as she can (too fast to count), hearers are left with an intuition of approximate number that tracks the actual number. Similarly, if someone asks you to make 50 taps on the table as fast as you can (too fast to count), the average number of taps you make across trials will approximate to 50. Likewise, if someone flashes up a set of dots on a screen (again, too briefly to count), you will have an impression of number that approximates the real number. Control experiments show that these are genuinely representations of abstract number, not volumes or surface-areas, for example. And although one might express what one experiences by saying something like, “There were about 50 dots on the screen”, there is widespread agreement that these numerosity-experiences do not embed exact-number concepts like *fifty*. (Otherwise they would be inaccessible to non-human animals and human infants.) The representation of number in question is more like a continuous curve, centered roughly on a magnitude that tracks, on average, the number 50.

Representations of numerosity are a component of, or at least accompany, one’s sensory experience. But they are not themselves sensory-specific. On the contrary, an approximate representation of 50 items can remain the same across auditory presentations, visual presentations, or mixed stimuli that combine both visual (dots on a screen) and auditory (tapping sound) components. So they are non-sensory. Moreover, they can surely be access-conscious. One often has a conscious impression of approximate number, which can ground a more-or-less accurate guess as to the exact number. The question is whether these numerosity-representations are also *phenomenally* conscious.

Our operationalization of phenomenal consciousness suggests that they are. Imagine a case where 50 dots have been flashed on a screen, giving one a rough impression of number. It seems one can coherently think, “Someone could be like me in all sensory respects, and could make guesses in cases like this that are just as accurate as mine, but who lacks *this* feeling [of approximate number]”. One can go on to think, “Such a person would be *merely* guessing (albeit reliably), without her guesses being grounded in *this* sort of impression of number.”

Likewise, it seems one can imagine a case parallel to that of color-deprived Mary. One can imagine someone whose numerosity faculty is inoperative, who nevertheless learns to count (perhaps with greater difficulty than normal), and goes on to become a world-famous cognitive neuroscientist who knows everything there is to know about approximate number processing. If suddenly cured of her deficit, and confronted with a novel collection of items, she might exclaim, “So *this* is what it is like to experience the numerosity of a set!”

We suggest, then, that not only valence but also approximate-number representations (when access-conscious) make a constitutive and irreducible contribution to one’s phenomenally-conscious experience. Since both are non-conceptual, as are sensory experiences of all kinds, it is beginning to look on inductive grounds as if *all* non-conceptual contents might be phenomenally conscious when access-conscious. We will consider one further case.

3.4. Temporality

Humans, like other animals, have a sense of time. This is underwritten by a variety of distinct temporal mechanisms, one of which tracks the approximate time of day, one of which operates on 24-h cycle governing wakefulness and sleep, and one of which is sensitive to intervals between events and collections of such intervals (Gallistel, 1990; Montemayor, 2012). It is widely agreed in the literature on episodic memory that a sense of some extent of *pastness* is a component of such memories, and is phenomenally conscious. There is said to be a distinctive experience of the form, *I experienced this so long ago*, where the embedded temporal representation is a *feeling* of elapsed time, not a conceptual representation of any sort (Michaelian, 2016; Tulving, 1983).⁹

Consider also interval memory. Suppose one is engaged in a task where, following a stimulus of some sort (a flash of light, say), one has a certain interval of time in which to do one thing (collect tokens at one location on a computer screen), after which one needs to do another (switch to another location; see Balci, Freestone, & Gallistel, 2009). Consider how it feels as one holds the cursor at the first location. One has a sense that *time is running out*, and an increasing feeling of urgency that one should switch to the other alternative. Put differently, one has the sense that the stimulus-event is receding into the past in a way that is approaching one’s target decision-point. These are non-conceptual experiences, which might *ground* a conceptual judgment like, “It is time to switch”, but which don’t seem to reduce to any combination of such judgments. They are also, plainly, non-sensory experiences. One’s impression of time can remain the same whether one is dealing with intervals between visual events or intervals between auditory ones.

⁹ It is often claimed that the contents of episodic memory are distinctively human – unavailable to be experienced even by animals capable of recalling what happened, when, and where – on the grounds that episodic memory requires an awareness of *oneself as experiencing* the events in question (Michaelian, 2016; Tulving, 1983). We are doubtful of this claim. While humans will *express* their episodic memories in metarepresentational terms, claiming to recall an *experience* of a certain sort (as such), we are doubtful that the content of the episodic memory is itself metarepresentational. Rather, it is a perspectival sensory-like representation of a worldly event of some kind, represented as having occurred a certain felt distance in the past. But this point is ancillary to our main focus here, which is specifically on the temporal component of episodic (and other) memories.

Moreover, it is obvious that these experiences are access-conscious. (One can report them, for example.) The question is whether they are also phenomenally conscious.

Our operationalization suggests that they are. It seems one can imagine someone who lacks a (conscious) sense of time, who can nevertheless behave as normal and make normal judgments of time. One can think, “Someone could be like me in all sensory respects and could perform just as successfully at this interval-tracking game, making accurate judgments of elapsed time, while lacking *this* sort of *feeling* of elapsed time.” Likewise, one can imagine Jackson-style cases. One can imagine someone congenitally lacking any sense of time (who has to do everything by the clock), who nevertheless learns all there is to know about the neural basis and functional organization of human beings’ temporal abilities. Once cured (perhaps by some future form of genetic engineering), she might exclaim when trying to judge the temporal interval between two events for the first time, “So *this* is what it feels like to sense that an event is receding into the past!”

There is a plausible case, then, for saying that our temporal experiences are non-conceptual and non-sensory, while making a constitutive and irreducible contribution to the phenomenal qualities of our experience. The inductive basis for claiming that *all* (as well as only) non-conceptual contents are phenomenally conscious when access-conscious is thereby significantly strengthened.¹⁰

3.5. Debate-lines redrawn

It seems that our proposed operationalization has proven its worth. Not only can it help resolve the debate over cognitive phenomenology in a principled and non-question-begging way, but it can also be used to evaluate the phenomenally-conscious status of a range of different kinds of non-conceptual experience that aren’t often considered in discussions of this sort (including valence, experiences of approximate number, and temporal experience).

In fact, the application of our operationalization suggests two further things. First, the present section has made a case that all non-conceptual states (when access-conscious) are irreducibly phenomenally conscious. And then since we argued in Section 2 that there is no irreducible phenomenology associated with thoughts or conceptual states – or in other words that the only mental states with irreducible phenomenology are *non*-conceptual – we now have reason to think that *all and only* non-conceptual states (when access-conscious) are irreducibly phenomenally conscious.

Second, our operationalization suggests that we should rethink the boundaries of the cognitive-phenomenology debate. As we just mentioned, we argued in Section 2 (as conservatives typically do) that conceptual states don’t have irreducible phenomenology. However, participants on both sides of the debate about cognitive phenomenology generally assume that for conservatives, the states with irreducible phenomenology, to which cognitive phenomenology is reduced, are all broadly sensory ones. Bayne and Montague (2011), for example, write that a conservative view “conceives of *sensory* phenomenology, broadly construed as above [to include images, moods, and the feelings associated with emotions], as the *only kind* of phenomenology there is” (2, emphasis ours). We take it that this correctly characterizes the standard conservative view. (For instance, see Tye & Wright, 2011, 329.) Note that they include within the domain of the sensory the aspects of emotion discussed in Section 3.1: the proprioceptive and interoceptive sensory experiences involved in facial expressions, bodily postures, “gut feelings”, and bodily arousal. But they make no mention of the valence component of emotional states, let alone of experiences of approximate number or time. Yet the present section has argued for the irreducible phenomenology of these types of content, which are non-conceptual *and yet not sensory* either.¹¹

As a result, we believe it is a mistake to characterize the competing views in the debate about cognitive phenomenology in terms of one’s stance on the reducibility or irreducibility of cognitive phenomenology to *sensory states*. For what should really be at issue is whether or not the phenomenology associated with thoughts and conceptual states is reducible to *other states with irreducible phenomenology* – whether those turn out to be broadly sensory or not. Someone might insist, of course, that the only people who deserve to be called “conservatives” are those who accept that irreducible phenomenology is broadly sensory; but it needs to be acknowledged that there is more than one way to *deny* such a claim. One can fail to be a conservative (in this sense) by being a typical liberal – by believing, that is, that thoughts and conceptual states also have irreducible phenomenology. Or one can fail to be a conservative (in this sense) by being a *non-conceptualist* – by believing, as we do, that broadly sensory states aren’t the *only* states with irreducible phenomenology (since valence, experiences of numerosity, and temporal feelings have irreducible phenomenology), while still denying that thoughts and conceptual states have irreducible phenomenology.

One advantage of taking our operationalization seriously, then, is that it can make salient forms of irreducible phenomenology that aren’t typically recognized in the context of this debate. Not only does this matter for the sake of completeness and clarity, but it also has substantive implications for the plausibility of rejecting the classic liberal position. Remember that phenomenal-contrast

¹⁰ Why are all such contents phenomenally conscious? We think this is because (a) reductive representationalism about phenomenal consciousness is correct, and (b) all and only such contents support the deployment of a special class of first-person concepts (often called “phenomenal concepts”), which lack any conceptual connections with the sorts of physical, functional, and intentional concepts that get held fixed in hard-problem-type thought experiments (Carruthers & Veillet, 2007). But these are not claims we propose to develop any further here, and the arguments of the present paper are independent of them.

¹¹ It might be wondered whether this claim remains true under *all* characterizations of “broadly sensory” states in the literature. (It is certainly true that valence, experiences of numerosity, and feelings of temporal passage don’t count as sensory given the way that *most* contributors to the debate have understood the notion.) For example, consider Chudnoff (2015a,b) who writes: “a mental state represents that p in a sensory way just in case it represents that p in a way that is dependent on current awareness of an environmental witness to p or a state that [like a visual image] is akin to such awareness” (2015b, 86). Given such a characterization, it might be wondered whether the types of content we have drawn attention to in this section like might qualify as sensory after all. The answer isn’t wholly clear, and Chudnoff doesn’t mention such cases. But our own feeling is that it would be quite a stretch to characterize of an experience of temporal distance as a “current awareness of an environmental witness to the past”. For what is it (beyond the feeling itself) that provides such witness? Similar points could be made about valence and experiences of numerosity.

arguments are inferences to the best explanation. Non-conceptualists like ourselves believe they can explain the sorts of contrasts these arguments highlight just as well as liberals by appealing to non-conceptual forms of phenomenology. It may then be important that the range of these non-conceptual kinds of phenomenology be fully acknowledged. If valence and time can be phenomenally conscious, for example, then the phenomenology of valence and time can get to figure in our explanation of the relevant contrasts. We will return to this point in Section 4.

Interestingly, Kriegel (2015) makes a related point in his discussion of emotional phenomenology. By the time he discusses this topic he takes himself to have established the irreducibility of cognitive phenomenology, conative phenomenology (the phenomenology of desiring or wanting), and of the phenomenology of entertaining a proposition (without being committed to its truth). He then sets out to defend a “conservative” account of emotional phenomenology, arguing that it is in fact reducible to other forms of phenomenology. But what motivates his reducibility claim with respect to emotional phenomenology is the detailed account he has provided of the other kinds of phenomenology to which emotional phenomenology could possibly be reduced. He writes, “Once we appreciate the richness and complexity of perceptual, cognitive, and conative aspects of emotional experience, I will suggest, it is hard to motivate positing an addition *sui generis* emotional phenomenology” (129). We believe, echoing Kriegel, that our operationalization provides a way to appreciate the richness, complexity, and range of forms of phenomenology that are neither sensory nor cognitive, to which non-conceptualists can appeal when reducing conceptual phenomenology itself.

4. Kriegel’s Zoe

We now discuss the example that is perhaps most challenging for our view: the case of Zoe. Like ourselves, Kriegel (2009, 2015) proposes to operationalize phenomenal consciousness as whatever gives rise to hard-problem thought experiments. He states the view as follows:

For any mental property F, F is a phenomenal property iff there is rationally warranted appearance (to a sufficiently reflective but otherwise normal actual subject) of a distinctive (e.g. empirical and nonderivative) explanatory gap between F and physical properties (2015, 52).

For Kriegel, as for Horgan (2011) and McClelland (2016), what would indicate the presence of the right type of explanatory gap is a certain kind of partial zombie. But Kriegel’s partial zombie is very different from Horgan’s and McClelland’s. Where Horgan gave Andy₃ the same causal/functional structure and sensory/affective phenomenology as Andy₁, Kriegel (2015) builds his partial zombie by *taking out* her sensory/emotional (non-cognitive) phenomenology. He imagines Zoe, who “enjoys no sensory phenomenology, no algedonic phenomenology [pleasure or pain] and no emotional phenomenology” (55). Moreover, he makes no effort to keep Zoe causally or functionally identical to an ordinary person, avoiding the kinds of worries we raised about the robust conceivability of Andy₃ in Section 2. Zoe lacks sensory, algedonic, and emotional phenomenology precisely because she is causally/functionally very different from ordinary people. This difference, in turn, may result from the fact that:

... the relevant parts of her brain are deformed, or are lesioned, or because – as they say in the genre – of a “cosmic incident”. Nonetheless, we can imagine this person having considerable underlying information-processing power, though one that does not translate into phenomenal experience (2015, 55).

Indeed, Zoe is described as a mathematical genius. As such, Kriegel imagines her undergoing “an episode of sudden realization of how a proof must go” (2015, 58). Zoe’s sudden realization “as well as the vivacity or intensity with which it is presented” are intuitively *phenomenal* states, Kriegel thinks. She would all of sudden *see* how the proof goes and *feel* the intensity of its presentation. This phenomenology, Kriegel argues, can’t possibly be reducible to sensory and/or emotional phenomenology since Zoe, by stipulation, doesn’t have any. Indeed, he maintains that his opponents can’t account for Zoe’s case because they believe that cognitive phenomenology is reducible to sensory and/or emotional phenomenology.

Note that Kriegel is not here claiming that the individual thoughts that constitute Zoe’s proof are phenomenally conscious; nor is he saying that what is phenomenally conscious is the thought that those propositions, when suitably chained or related, constitute a proof. If he were, his argument would be vulnerable in just the way that claims of cognitive phenomenology generally are, as we discussed in Section 2. For it is far from obvious that once the access-conscious status of a thought is fixed there is anything specific to that individual thought remaining, which could be imagined away (or in the present case, imagined as residue) in a partial-zombie thought experiment. Instead, what Kriegel emphasizes is the event of *seeing* how the proof should go, experienced with a certain intensity.

We should point out that Kriegel has presented us with a delicate and challenging task. In order to imagine Zoe’s mental life one has to start from a case of sudden realization with which one is introspectively familiar. One then has to subtract out all sensory and affective experience to focus on what remains. Kriegel *thinks* that when one does this there is something remaining – a *feeling of seeing* – that is phenomenally conscious but neither reduces to the propositions/thoughts in question nor is sensory/affective in nature. But how can he be sure? How can one tell, just by introspecting one’s own experience, that removing all sensory and affective components wouldn’t take away the feeling of seeing as well? To assume, as Kriegel seems to, that one is in a position to answer these questions introspectively is evidence of the pervasiveness of the assumption that relationships among kinds of phenomenology will be introspectively transparent and obvious. But as we pointed out in Section 1, there is no reason to think that introspection could settle the question whether or not one type of phenomenology is reducible to another. One should thus be quite suspicious of claims about what it would be like for a person (in this case Zoe) if we removed large portions of her mental life.

Our discussion in Section 3 provides materials for an additional critique of Kriegel’s characterization of Zoe, however. For there

are plausibly two aspects to a feeling of seeing (in addition to the thought-contents in question, of course). Each of these admits of degrees, thus explaining how the feeling of seeing can be more or less intense. The first is what psychologists describe as a feeling of *fluency* (Oppenheimer, 2008; Topolinski & Reber, 2010). When one is struggling to see how a set of propositions might constitute a proof, processing of those thoughts and the relations between them is *disfluent*. When one comes to see the solution, in contrast, processing suddenly becomes *fluent*.

What does this mean? In part, fluency is a temporal phenomenon. It means that one can move through the thoughts that constitute the proof swiftly. But we have already argued, in Section 3.4, that representations of temporal passage and temporal intervals are phenomenally conscious. So there *would* be a component of Zoe’s mental life that would be neither sensory/affective nor conceptual, but which would still be phenomenally conscious. Fluency is also an affective phenomenon, however. Disfluent processing is *effortful* processing, whereas fluent processing is effortless. The best account of such feelings of cognitive effort is that they are constituted by negative valence directed at one’s current activity (Kurzban, Duckworth, Kable, & Myers, 2013). So an event of *seeing* the proof will be accompanied by loss of the negative valence previously directed at one’s activity of thinking the thoughts in question. But here, too, we have argued (in Section 3.2) that valence is neither sensory nor conceptual, but is nevertheless phenomenally conscious. So while Kriegel is right that this aspect of Zoe’s mental life would be phenomenally conscious, this does nothing to support the existence of *conceptual* phenomenology.

In addition to fluency, the other main component of a feeling of seeing is likely to be a feeling of *confidence*. When one comes to see a proof for the first time one doesn’t merely entertain the thought that it is a proof, one *feels confident* that it is. Again, the main contributor to feelings of confidence is valence (in this case positive valence). To have confidence in one’s answer to a question is for that answer to *seem gradedly-good to one* (that is, for it to be positively valenced). And likewise, to have confidence in the claim that a certain sequence of propositions amounts to a proof is for that claim to *seem gradedly-good to one*. It is plausible, then, that Zoe *would* have phenomenal states in the absence of sensory ones when she sees a new proof. But this fails to support the existence of cognitive (conceptual) phenomenology. It is, on the other hand, consistent with the suggestion that all and only non-conceptual states (when globally accessible) are phenomenally conscious.

A potential problem for this treatment of the case of Zoe is that Kriegel (2015) stipulates, not just that Zoe is lacking all sensory and affective (emotional) forms of experience, but that she is also lacking in *algedonic* experience (feelings of pleasure and pain). However, it is unclear whether this implies, for Kriegel, that Zoe lacks all forms of (non-sensory) positive and negative valence. For it may be that he has in mind properties that would be specific to proprioceptive and interoceptive forms of bodily experience, rather than valence more generally. But either way, we believe our point stands.

If, as most people do, Kriegel understands pleasure and pain narrowly, to encompass only sensory pleasure and sensory pain, then Zoe’s lack of algedonic phenomenology won’t mean that she lacks the phenomenology of valence. For we claim that the evaluative component in these sensory states (positive and negative valence) is the same as that involved in affective states generally, including those that are much more abstract in nature, such as the positive valence that can attach to the claim that something is a proof.

But even if Kriegel *does* intend “algedonic” to be taken more broadly, we think that in discussions of phenomenal consciousness it is easy to overlook feelings of valence that are untethered to any form of sensory presentation. Indeed, it is easy to overlook the phenomenal contribution made by the valence component of even sensory states like physical pain, which is one reason why it was worthwhile to argue the point, as we did in Section 3.2. So Kriegel may simply have overlooked the presence of valence phenomenology when introspecting his own feelings of insight and engaging in the subtractive exercise required to envisage Zoe. One can be introspectively familiar with feelings of insight without realizing how those feelings parse into a number of components (including a release from mental effort, swift passage among the component propositions, and positive valence directed at the claim that those thoughts constitute a proof).

Some have raised worries about explanations of the relevant contrast in terms of the kinds of phenomenology we discuss here, like fluency or confidence. Chudnoff (2015a,b), for instance, considers the possibility that a situation like Zoe’s (the sudden understanding of a mathematical proof) could be adequately explained by “a relief in general tension”, or a “general feeling of getting it” or a “feeling of self-satisfaction” (2015b, 100). To rebut this, he asks us to imagine that “a pill that you took earlier kicks in and relieves your general tension. Now you feel the relief of general tension but still don’t ‘see’ that the proposition is true” (Chudnoff, *ibid*). But this move is too quick, as Chudnoff himself notes. The claim we make here is not that Zoe experiences a relief in *general tension*, for that is arguably a sensory (proprioceptive and/or interoceptive) experience. Nor does she experience a sudden but general feeling of self-confidence. On the contrary, the valence component of confidence has to be *directed* adequately. Indeed, it has to be directed at the content of a mathematical thought.¹²

It might also be objected that if a phenomenal component of one’s experience (its valence) is *directed* at the content of a thought, then that thought must make a constitutive contribution to one’s overall phenomenal state. But *directedness* can be cashed out in functional terms, in the roles that the propositional and valence components respectively play in decision making (with the valence component motivating one to *assert* the thought component, for example, or motivating one to attempt to fix it in long-term memory for later use). And the objection again seems to lose sight of the point developed in Section 2, that conscious states can often parse into components, some of which are both access-conscious and phenomenally-conscious (here, valence), but others of which are

¹² Chudnoff adds that this move – insisting that the feelings in question be adequately directed – is still less than ideal, because it fails to make clear the similarities between suddenly seeing a mathematical truth (as opposed to merely entertaining the thought that it might be true) and suddenly *seeing the mail in your mailbox* (as opposed to merely entertaining the thought that it might be there). But it is far from obvious why these two acts of “suddenly seeing” should be given similar explanations.

merely access-conscious (here, the content of the thought).

We conclude, then, that the case of Zoe does nothing to support the existence of irreducible conceptual phenomenology. What remains phenomenally conscious in her mental life can be adequately accounted for in terms of features of her mentality that are neither sensory nor conceptual (including positive and negative valence as well as her sense of time and temporal passage).

5. Concluding discussion

We continue to believe that the way to address questions about the scope of phenomenal consciousness is by considering whether hard-problem-type thought experiments are applicable. But we have accepted McClelland's (2016) point that only variants of the zombie thought experiment are a reliable diagnostic tool for the *absence* of phenomenal consciousness. Thus employed, we think that the case for conceptual or propositional forms of irreducible phenomenology is weak. But we have also argued for the existence of a variety of forms of phenomenology that are both non-conceptual and non-sensory, including valence, a sense of approximate number, and feelings of temporal extent and passage. Thus we no longer think that phenomenal consciousness is restricted to sensory states, no matter how broadly characterized. Rather, all non-conceptual states (whether sensory or non-sensory) are phenomenally conscious when access-conscious.

The upshot can be thought of as a weakened form of conservatism, since we now allow that non-sensory, amodal, states can make an intrinsic contribution to the phenomenal properties of our lives. But others might claim victory for a weakened form of liberalism, since the amodal representations in question can be thought of as broadly cognitive (albeit non-conceptual) in character. It is better, we think, to drop the labels altogether. Our claim is that all and only access-conscious non-conceptual contents are phenomenally conscious. To oppose us, it needs to be argued that concepts and/or propositions can make a constitutive, irreducible, contribution to the phenomenally-conscious properties of a mental state.

A final upshot of the paper is that there is a real (and not merely conceptual) distinction between access-consciousness and phenomenal consciousness. But the distinction doesn't apply at the level of whole states (as Block, 2011, claims), but rather to components of states. Specifically, many types of access-conscious state – like hearing someone assert that ducks are cute – can be parsed into components that are both access-conscious and phenomenally conscious, on the one hand (this will be true of all non-conceptual contents of the state, sensory or not), and components that are *merely* access-conscious, on the other (in this case the proposition, *ducks are cute*).

Acknowledgments

We are grateful to Timothy Lane for his advice, and to three anonymous referees for their insightful (and sometimes challenging) comments on an earlier version of this article.

References

- Armstrong, D. (1962). *Bodily sensations*. London: Routledge.
- Bain, D. (2014). What makes pains unpleasant? *Philosophical Studies*, 166, S69–S89.
- Balci, F., Freestone, D., & Gallistel, C. R. (2009). Risk assessment in man and mouse. *Proceedings of the National Academy of Sciences*, 106, 2459–2463.
- Barth, H., Kanwisher, N., & Spelke, E. (2003). The construction of large number representations in adults. *Cognition*, 86, 201–221.
- Bartra, O., McGuire, J., & Kable, J. (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, 76, 412–427.
- Bayne, T. (2009). Perceptual experience and the reach of phenomenal content. *Philosophical Quarterly*, 59, 385–404.
- Bayne, T., & Montague, M. (Eds.). (2011). *Cognitive phenomenology*. Oxford: Oxford University Press.
- Bechara, A., Damasio, H., & Damasio, A. (2000). Emotion, decision making and the orbitofrontal cortex. *Cerebral Cortex*, 10, 295–307.
- Bermúdez, J. (2003). *Thinking without words*. Oxford: Oxford University Press.
- Bermúdez, J. (2015). Nonconceptual mental content. In E. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/archives/fall2015/entries/content-nonconceptual>.
- Block, N. (1986). Advertisement for a semantics for psychology. *Midwest Studies in Philosophy*, 10, 615–678.
- Block, N. (1995). A confusion about the function of consciousness. *Behavioral and Brain Sciences*, 18, 227–247.
- Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30, 481–499.
- Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Science*, 12, 567–575.
- Block, N. (2014). Rich conscious perception outside focal attention. *Trends in Cognitive Sciences*, 18, 445–447.
- Botterill, G., & Carruthers, P. (1999). *The philosophy of psychology*. NY: Cambridge University Press.
- Carruthers, P. (2000). *Phenomenal consciousness*. NY: Cambridge University Press.
- Carruthers, P. (2015b). *The centered mind: What the science of working memory shows us about the nature of human thought*. Oxford: Oxford University Press.
- Carruthers, P. (2015a). Block's overflow argument. *Pacific Philosophical Quarterly*. <http://dx.doi.org/10.1111/papq.12152> Online first.
- Carruthers, P. (2017). Valence and value. *Philosophy and Phenomenological Research*. <http://dx.doi.org/10.1111/phpr.12395> Online first.
- Carruthers, P. (2009). Invertebrate concepts confront the generality constraint (and win). In R. Lurz (Ed.), *The philosophy of animal minds: New essays on animal thought and consciousness* (pp. 89–107). NY: Cambridge University Press.
- Carruthers, P., & Veillet, B. (2007). The phenomenal concept strategy. *Journal of Consciousness Studies*, 14, 212–236.
- Carruthers, P., & Veillet, B. (2011). The case against cognitive phenomenology. In T. Bayne, & M. Montague (Eds.), *Cognitive phenomenology* (pp. 35–56). Oxford: Oxford University Press.
- Chalmers, D. (1996). *The conscious mind*. Oxford: Oxford University Press.
- Chudnoff, E. (2015b). Phenomenal contrast arguments for cognitive phenomenology. *Philosophy and Phenomenological Research*, 90, 82–104.
- Chudnoff, E. (2015a). *Cognitive phenomenology*. London: Routledge.
- Cohen, M., & Dennett, D. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, 15, 358–364.
- Damasio, A. (1994). *Descartes' error*. London: Papermac.
- Dehaene, S. (2014). *Consciousness and the brain*. London: Penguin Press.

- Dretske, F. (1988). *Explaining behavior*. Cambridge, MA: MIT Press.
- Dummett, M. (1994). *Origins of analytical philosophy*. Cambridge, MA: Harvard University Press.
- Durso, G., Luttrell, A., & Way, B. (2015). Over-the-counter relief from pains and pleasures alike: Acetaminophen blunts evaluation sensitivity to both negative and positive stimuli. *Psychological Science*, 26, 750–758.
- Ellsworth, P., & Scherer, K. (2003). Appraisal processes in emotion. In R. Davidson, K. Scherer, & H. Goldsmith (Eds.), *Handbook of affective sciences*. New York: Oxford University Press.
- Evans, G. (1982). *The varieties of reference*. Oxford: Oxford University Press.
- FitzGerald, T., Seymour, B., & Dolan, R. (2009). The role of human orbitofrontal cortex in value comparison for incommensurable objects. *Journal of Neuroscience*, 29, 8388–8395.
- Fodor, J. (1990). *A theory of content and other essays*. Cambridge, MA: MIT Press.
- Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- Gallistel, C. R., Mark, T., King, A., & Lantham, P. (2001). The rat approximates an ideal detector of rates of reward. *Journal of Experimental Psychology: Animal Behavior Processes*, 27, 354–372.
- Horgan, T. (2011). From agentive phenomenology to cognitive phenomenology: A guide for the perplexed. In T. Bayne, & M. Montague (Eds.), *Cognitive phenomenology* (pp. 57–78). Oxford: Oxford University Press.
- Izard, V., Sann, C., Spelke, E., & Streri, A. (2009). Newborn infants perceive abstract numbers. *Proceedings of the National Academy of Sciences*, 106, 10382–10385.
- Jackendoff, R. (1987). *Consciousness and the computational mind*. Cambridge, MA: MIT Press.
- Jackendoff, R. (2012). *A user's guide to thought and meaning*. Oxford: Oxford University Press.
- Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32, 127–136.
- Jackson, F. (1986). What Mary didn't know. *Journal of Philosophy*, 83, 291–295.
- Jordan, K., MacLean, E., & Brannon, E. (2008). Monkeys match and tally quantities across senses. *Cognition*, 108, 617–625.
- Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, 14, 301–307.
- Kriegel, U. (2009). *Subjective consciousness*. Oxford: Oxford University Press.
- Kriegel, U. (2015). *The varieties of consciousness*. Oxford: Oxford University Press.
- Kurzban, R., Duckworth, A., Kable, J., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *Behavioral and Brain Sciences*, 36, 661–679.
- Leknes, S., & Tracey, I. (2008). A common neurobiology for pain and pleasure. *Nature Reviews Neuroscience*, 9, 314–320.
- Levy, D., & Glimcher, P. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology*, 22, 1027–1038.
- Lieberman, M. (2013). *Social*. New York: Crown Publishers.
- Lieberman, M., & Eisenberger, N. (2009). Pains and pleasures of social life. *Science*, 323, 890–891.
- McClelland, T. (2016). Gappiness and the case for liberalism about phenomenal properties. *Philosophical Quarterly*, 66, 536–558.
- McDowell, J. (1994). *Mind and world*. Cambridge, MA: Harvard University Press.
- McGinn, C. (1982). The structure of content. In A. Woodfield (Ed.), *Thought and object*. Oxford: Oxford University Press.
- Michaëlian, K. (2016). *Mental time travel*. Cambridge, MA: MIT Press.
- Millikan, R. (1984). *Language, thought, and other biological categories*. MIT Press.
- Montemayor, C. (2012). *Minding time*. Leiden: Brill.
- Oppenheimer, D. (2008). The secret life of fluency. *Trends in Cognitive Sciences*, 12, 237–241.
- Papineau, D. (1987). *Reality and representation*. Oxford: Blackwell.
- Peacocke, C. (1992). *A study of concepts*. Cambridge, MA: MIT Press.
- Phillips, I. (2011). Perception and iconic memory: What Sperling doesn't show. *Mind & Language*, 26, 281–311.
- Pitt, D. (2004). The phenomenology of cognition or What is it like to think that P? *Philosophy and Phenomenological Research*, 69, 1–36.
- Plassmann, H., O'Doherty, J., & Rangel, A. (2010). Appetitive and aversive goal values are encoded in medial orbitofrontal cortex at the time of decision making. *Journal of Neuroscience*, 30, 10799–10808.
- Prinz, J. (2004). *Gut reactions*. Oxford: Oxford University Press.
- Prinz, J. (2012). *The conscious brain*. Oxford: Oxford University Press.
- Ruff, C., & Fehr, E. (2014). The neurobiology of rewards and values in social decision-making. *Nature Reviews Neuroscience*, 15, 549–562.
- Russell, J. (2003). Core affect and the psychological construction of emotion. *Psychological Review*, 110, 145–172.
- Siewert, C. (1998). *The significance of consciousness*. Princeton, NJ: Princeton University Press.
- Siewert, C. (2011). Phenomenal thought. In T. Bayne, & M. Montague (Eds.), *Cognitive phenomenology* (pp. 236–267). NY: Oxford University Press.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, 74, 1–29.
- Stazicker, J. (2011). Attention, visual consciousness and indeterminacy. *Mind & Language*, 26, 156–184.
- Strawson, G. (1994). *Mental reality*. Cambridge, MA: MIT Press.
- Strawson, G. (2011). Cognitive phenomenology: Real life. In T. Bayne, & M. Montague (Eds.), *Cognitive phenomenology* (pp. 285–325). Oxford: Oxford University Press.
- Topolinski, S., & Reber, R. (2010). Gaining insight into the “Aha” experience. *Current Directions in Psychological Science*, 19, 402–405.
- Tulving, E. (1983). *Elements of episodic memory*. Oxford: Oxford University Press.
- Tye, M. (1995). *Ten problems of consciousness*. Cambridge, MA: MIT Press.
- Tye, M. (2000). *Consciousness, color, and content*. Cambridge, MA: MIT Press.
- Tye, M., & Wright, B. (2011). Is there a phenomenology of thought? In T. Bayne, & M. Montague (Eds.), *Cognitive phenomenology* (pp. 326–344). Oxford: Oxford University Press.
- Wilkinson, H., Davidson, K., & Davidson, R. (1999). Bilateral anterior cingulotomy for chronic noncancer pain. *Neurosurgery*, 45, 1129–1136.