




Phenomenal intentionality: reductionism vs. primitivism

Philip Woodward 

Department of Philosophy, Valparaiso University, Valparaiso, IN, USA

ABSTRACT

This paper explores the relationship between phenomenal properties and intentional properties. In recent years a number of philosophers have argued that intentional properties are sometimes necessitated by phenomenal properties, but have not explained why or how. Exceptions can be found in the work of Katalin Farkas and Farid Masrour, who develop versions of reductionism regarding phenomenally-necessitated intentionality (or 'phenomenal intentionality'). I raise two objections to reductive theories of the sort they develop. Then I propose a version of primitivism regarding phenomenal intentionality. I argue that primitivism avoids the pitfalls of reductionism while promising broad explanatory payoffs.

ARTICLE HISTORY Received 23 May 2017; Accepted 3 April 2018

KEYWORDS Phenomenal properties; intentional properties; phenomenal intentionality

1. Introduction

Consciousness and content are the two traditional 'marks of the mental.' Consciousness is the experiential quality of mental states – what it is like for the subject to be in them. Content is the 'directedness' of mental states – whatever mental states are *about* or *represent*. Whether a mental state is conscious, and the particular way in which it is conscious, is a matter of the *phenomenal properties* that it includes. Whether a mental state has content, and the particular content it has, is a matter of the *intentional properties* it includes.

In recent years, a number of philosophers have argued that phenomenal and intentional properties can have a very close relationship: some conscious mental states are such that their phenomenal properties *necessitate* certain intentional properties. I'll call phenomenally-necessitated intentionality 'phenomenal intentionality.'¹ Its defenders have argued for it by making an observation about

either the phenomenology or intentionality of conscious mental states and then inferring the existence of phenomenal intentionality as the best explanation of the observation. These arguments have tended to fall into three types. (1) *Introspective arguments*. Observation: some conscious states are such that a subject can discern their accuracy conditions from the first-person perspective, and/or discriminate their contents from one another from the first-person perspective. For example, the conscious visual state ordinarily produced when staring at a picture hanging on a wall is accurate only if in fact a picture hangs from a wall at some distance from the subject, and all of this is introspectively manifest.² Or again, on the basis of their visual experiences, subjects can distinguish visually presented objects from one another and from the visual background.³ (2) *Phenomenal duplication arguments*. Observation: some pairs of conscious states are such that they can share intentional contents despite dramatic differences in the external situation of the subjects who experience them, so long as the two states are sufficiently similar phenomenally. For example, the phenomenal state ordinarily produced when looking at a lemon presents lemon-directed intentional content, whether or not the subject is hallucinating, or even whether the subject is a brain in a vat.⁴ (3) *Intentional contrast arguments*. Observation: some conscious states are such that shifts in intentional contents cannot occur unaccompanied by shifts in phenomenal character. For example, if a subject initially seems to see a wall containing a concave niche, and then – after realizing she is looking at a tromp l’oeil painting – seems to see a flat wall, the change in represented spatial features is accompanied by a change in what her experience is like.⁵ According to defenders of phenomenal intentionality, the best explanation of each of these observations is that there is a suite of intentional properties that supervene on the phenomenal properties of the states in question.

All of the examples I have invoked so far have been perceptual. Some philosophers have gone farther, offering versions of all three arguments that are meant to apply to conscious *cognitive* states. These arguments have been met with much more resistance than those restricted to perceptual cases.⁶ For present purposes, I am going to assume that phenomenal intentionality does exist, if only in the domain of conscious perception, and that the abductive arguments given in its favor are jointly compelling. My interest is in what could explain phenomenal intentionality: why is there sometimes this modally strong relationship between a mental states phenomenal and intentional properties? My answer is a form of *primitivism* regarding phenomenal intentionality: there is a *sui generis* type of phenomenology that is intrinsically intentional. Primitivism stands in contrast to *reductionism* regarding phenomenal intentionality, according to which there is nothing more to phenomenal intentionality than suitably structured, garden-variety phenomenology. The primary aim of my paper is to argue in favor of primitivism over reductionism.

In order to get clearer on the phenomenon in need of explaining, it will be helpful to distinguish between four different types of intentional property that

advocates of phenomenal intentionality might claim supervene on phenomenal properties:

- (1) Phenomenal representation: the property of (apparently) being about something or other.
- (2) Phenomenal content: the property of (apparently) being about particular intentional contents (at some level of specificity).
- (3) Phenomenal objectivity: the property of (apparently) being about some concrete, mind-independent object or other.
- (4) Phenomenal reference: the property of being about particular mind-independent objects.

To illustrate: suppose that Elena is thinking about her black lab, Rocky (and that she is thinking of him *as* black and *as* a dog). Suppose, further, that a believer in phenomenal intentionality claimed that this intentional property (thinking about Rocky) is necessitated by the phenomenal properties Elena is instantiating. If this claim is accurate, then Elena's mental state exhibits phenomenal representation in virtue of representing *something*; phenomenal content, in virtue of representing Rocky's *blackness*; phenomenal objectivity, in virtue of representing something *as a dog*, i.e., as a mind-independent object; phenomenal reference, in virtue of representing *Rocky*. How and whether these features hang together in various mental states will depend, to a certain extent, on contested matters within philosophical semantics and philosophy of mind.⁷

My focus in what follows will be on *phenomenal content*. I am thus assuming that phenomenal intentionality involves, at a minimum, the supervenience of particular intentional contents on the phenomenal character of some mental states. In what follows, I will use the term 'phenomenal intentionality' to denote phenomenal content unless I specify otherwise.

In Sections 2–4, I use the proposals of Katalin Farkas and Farid Masrour as a jumping-off point for reductionism, and then I raise two objections: first, their strategy cannot be applied to imaginative phenomenal intentionality; second, their strategy does not provide a perspicuous reduction even of perceptual phenomenal intentionality.

In Section 5, I defend my own proposal. Primitivism is a hard sell: we should only endorse it if prospects for reductionism look truly dismal, and if primitivism promises rich explanatory rewards in its own right. I argue that this is indeed the situation. Phenomenal-intentional properties have a crucial but limited role to play, not only in a theory of phenomenal intentionality, but in broader epistemological and semantic theorizing as well.

2. Reductionism

Broadly speaking, there are two ways of explaining the existence of phenomenal intentionality: (a) phenomenal intentionality is a primitive feature of conscious

states, over and above familiar, non-intentional phenomenological features; (b) phenomenal intentionality can be reduced to suitably-structured, garden-variety phenomenology. I take it that the latter option is preferable, all else being equal, because more parsimonious – it does not require expanding our inventory of the basic types of phenomenal property.⁸

What would a reduction of phenomenal intentionality look like? Consider, by way of an analogy, what it is for some feature of a macro-level object – say, the climbability of a ladder – to reduce to more basic features of that object. The rails and rungs of a ladder are not themselves climbable. But once we understand the architecture of a ladder, we are in a position to see that the features of these parts (their shape and rigidity and so on), when those parts are suitable arranged, give rise to a climbable composite. Similarly, supplying a reduction of phenomenal intentionality requires showing how the suitable arrangement of non-phenomenal-intentional components of a phenomenal state gives rise to a state that exhibits phenomenal intentionality. A reductionist project along these lines has recently been explored by Katalin Farkas and Farid Masrour.

Katalin Farkas (2013) observes that the difference between phenomenal episodes that seem to present an objective world and those that do not runs broadly parallel with the difference between phenomenal episodes that admit of certain systematic correlations between phenomenal properties and those that do not. Auditory and visual sensations, for example, are often systematically correlated: sounds grow fainter as our visual image of the sound-source grows smaller. Further, experiences of these sorts systematically correlate with agential and proprioceptive experiences related to bodily movement, as we move and act and investigate. And it is precisely experiences such as these – experiences which exhibit such systematic, inter-modal correlations – that seem to present an objective world. In contrast, pain-sensations and visual ephemera such as after-images are not subject to these same sorts of inter-modal correlations, and they do *not* seem to present an objective world. She writes:

[A] difference between ordinary perceptual experiences on the one hand, and experiences that don't seem to present experience-independent objects – after-images, phosphenes – on the other...[is] that the simpler phenomenal features of perceptual experiences are organized into a systematic, cross-modally coherent and predictable order. This order is what I call the 'structure' of the experience.⁹

Noting this close connection between phenomenal 'structure' on the one hand and phenomenal intentionality on the other, Farkas proceeds to *reduce* the latter to the former:

A feeling may just be a feeling and not present or represent anything. However, when these sensory features are received by the subject in a highly organized and predictable structure...the experience may become suggestive of the presence of something *beyond* this experience, namely, an experience-independent object. Perceptual intentionality is thus constituted by the structure of sensory phenomenal features: by the way these features hang together and respond to movement and inquiry.¹⁰

Let's call phenomenal properties that are decidedly not intrinsically intentional 'somatosensory properties.' Somatosensory properties comprise visual, auditory, tactile, proprioceptive and affective sensations.¹¹ Farkas's proposal is that sensory experiences seem to present mind-independent states of affairs when and only when somatosensory properties, both intra-modally and inter-modally, covary in certain ways, and hence that phenomenal intentionality ought to be explained in terms of such covariance.

If that were the sum total of Farkas's proposal, it would not get us very far. Depending on what sorts of experiential features can contribute to the relevant patterns, the proposal either under-predicts or over-predicts intentional contents. Consider a case of fairly static phenomenal goings-on: waking up in the morning and being phenomenally presented with the appearance of your bedroom's white ceiling at some distance above you. In such a case there is very little 'movement and inquiry' for your visual phenomenology to covary with. By hypothesis, your experience does present intentional contents to you (you seem to see a white surface at some distance), but because your experience is so sparse and static, it is hard to locate much phenomenal covariance within it. In response, Farkas could say that your experience exhibits covariance between *absences* – viz., between your absence of movement, on the one hand, and an absence of a change in visual phenomenology, on the other. But if absences can contribute to the relevant patterns, all experiential episodes will be too richly patterned to explain where intentional content comes from. Add to our example that your first few moments upon waking include a steady headache. The covariance between your absence of movement and the absence of a change in head-pain does not make it seem to you that there is an object external to the mind that is painful.

So, something needs to be added. Farkas hints at one possibility: let phenomenal patterns accrue over one's lifespan.¹² Thus what fixes the phenomenal-intentional content of a subject's phenomenal state at a time is a matter of the patterns the somatosensory properties included in that state *have tended to be part of* in the past. While this move may go some distance in helping explain the ceiling-staring case – perhaps your ceiling-caused visual phenomenology *does* tend to covary with movement and inquiry, if we are taking into consideration all past ceiling-gazing episodes – puzzles will remain. Perhaps whitish color fields have tended to be correlated in the past with agential phenomenology associated with getting out of bed; but perhaps they have also been correlated with agential phenomenology associated with crossing a snowy meadow. Which pattern is constitutive of the phenomenal-intentional content you enjoy immediately upon waking? Suppose that you are gazing at the ceiling from a slightly different angle than ever before. Is such an unprecedented phenomenal state intrinsically contentless, or is its content supplied by those patterns that phenomenology *sufficiently similar to it* has tended to be part of? In short: some

additional element is needed to specify the types of patterns that are relevant to determining content.

A proposal similar to Farkas's but that seems to build in this extra element from the start is that of Farid Masrou (2013). Now, Masrou's explanatory target is phenomenal objectivity rather than phenomenal content. I am not certain whether he believes in phenomenal content, or if he does, whether he thinks that the same resources he uses to account for phenomenal objectivity can likewise account for phenomenal content.¹³ Regardless, it will be helpful to bring those resources to bear, especially because they seem to be the right resources to use to develop reductionism.

Masrou says that, in addition to phenomenal covariance, the constitution of phenomenal objectivity involves a special form of phenomenology that tracks the relevant patterns, an idea inspired by his reading of Kant. Masrou calls the resulting constellation of phenomenal elements 'schematic dynamical unity,' and explains it as follows:

In the Kantian view, in contrast [to a simple regularity account], the experience of schematic dynamical unity requires the activation of schematic representations which are associated with a specific form of phenomenology. A schematic representation is in effect a detector whose job is to test whether certain dynamic relations between the values of certain representations obtain. An abstract arbitrary rule that does not correspond to such a detector does not count as a schema.¹⁴

Imagine walking toward a tree as you are looking at it. When you get closer, the visual angle through which you see the tree grows in size in an inverse relation to your distance from the tree... There is something that it is like to experience the visual angle and the relative distance as covarying in the particular law-like manner that they do and this additional phenomenological element is over and above the law-like covariance.¹⁵

The general shape of Masrou's view (as I read him) is as follows. Not only are our sensations systematically correlated in certain ways. We also have phenomenological awareness of a subset of these correlations, as such. Relevant types of phenomenal correlation (*schemas*) are tracked by a dedicated psychological mechanism, outputs of which are 'schematic representations.' Schematic representations show up in consciousness as phenomenal highlightings of the obtaining of correlations. In particular, this highlighting comes in the form of what Masrou calls 'implicit anticipatory conditionals' – i.e. a sense of what's expected and what's not with respect to the unfolding of our perceptual experiences. When our experiences conform to these implicit anticipatory conditionals, and *because* they do so, they seem to present us with an objective world. The right sort of phenomenal awareness of the right sort of phenomenal patterns constitutes phenomenal objectivity.

I'll call Farkas's view, augmented by Masrou's proposals, a 'covariance theory' of phenomenal intentionality. We can distinguish between two types of covariance theory: restricted versions, intended to explain phenomenal content in a single psychological domain (e.g. the domain of the perceptual), and expansive

versions, intended to explain phenomenal content in multiple psychological domains. (Note that Farkas explicitly restricts her proposal to the perceptual domain.¹⁶) In the next two sections, I levy two objections against the covariance theory. The target of the first is a version of the covariance theory intended to explain phenomenal intentionality in the *imagination* – and, I argue, any version of the theory designed to explain perceptual phenomenal intentionality ought to be expansive enough to apply to the imagination. The target of the second objection is any covariance theory whatever, including versions restricted to the perceptual domain. As we'll see, the problems I mention in connection with my second objection are liable to plague any reductive theory of phenomenal intentionality, whether or not it is a version of the covariance theory.

3. First objection to reductionism: inapplicability to the imagination

My first objection has to do with the facts that (a) phenomenal intentionality can be found outside of the realm of the perceptual, viz., in imagination, but (b) it makes no sense to explain *imaginative* phenomenal intentionality in terms of phenomenal covariance. I'll present my objection in the form of an argument against the covariance theory, and then discuss each premise.

- Premise 1: Some imaginative states have phenomenal intentionality.
- Premise 2: The phenomenal intentionality of imaginative states cannot be explained in terms of phenomenal patterns or anticipations of phenomenal patterns.
- Premise 3: The phenomenal intentionality of imaginative states is to be explained in the same fashion as that of perceptual states.
- Conclusion: The covariance theory of the phenomenal intentionality of perceptual states is false.

Premise 1: Some imaginative states have phenomenal intentionality. My defense of this is that versions of each of the three arguments discussed above in favor of perceptual phenomenal intentionality can be framed to support imaginative phenomenal intentionality. Are imaginative states such that subjects can discriminate their contents from one another from the first-person perspective (per the introspective argument)? Clearly: I have no trouble discriminating an image as of a red cube from an image as of a red sphere, for example. Are pairs of imaginative states such they can share intentional contents despite dramatic differences in the external situation of the subjects who experience them, so long as the two states are sufficiently similar phenomenally (per the phenomenal duplication argument)? Clearly: the contents of imaginative states are even more obviously independent of the subject's external situation than are the contents of perceptual states, as I can imagine a red sphere no matter which environment I happen to be in.¹⁷ Are imaginative states such that aspects of their intentional contents cannot occur unaccompanied by shifts in phenomenal

character (per the intentional contrast argument)? Clearly: a shift in imagining a convex Necker cube to a concave Necker cube is a shift in what my imaginative state is like for me. Those who are moved by similar considerations to grant the existence of *perceptual* phenomenal intentionality should just as readily grant the existence of phenomenal intentionality in the domain of the imagination.

Premise 2: The phenomenal intentionality of imaginative states cannot be explained in terms of phenomenal patterns or anticipations of phenomenal patterns. According to the covariance theory, phenomenal states come to have a specific intentional content when those states exhibit certain phenomenal patterns characteristic of that content, or when the subject anticipates that such patterns would obtain. It does seem right to say, of conscious perceptual states, that specific contents are accompanied by phenomenal patterns of a certain sort, or at least by an anticipation of such patterns. For example: When I perceive a red rubber ball, I anticipate that the visual experience of the ball will take up a larger portion of my visual field as I experience walking toward the ball, that the ball will appear spherical unless a contiguous object appears to squish it, that it will retain its color so long as illumination conditions appear constant, that its visual appearance will be accompanied by certain tactile qualities if I experience reaching out and touching it, and so forth. But in the domain of the imagination, phenomenal patterns can completely break down. I can imagine a ball whose visual appearance shrinks in size as I move toward it; that de-forms spontaneously; that changes color willy-nilly; that doesn't feel like anything to touch. And it is not as though such imagined behavior is a surprising breakdown of what I anticipated. Because I know that my imaginative states are under the control of my will, I need not anticipate that they will unfold in any particular way. It is thus possible for my imaginative state to exhibit phenomenal intentionality without exhibiting any corresponding phenomenal patterns, and without my anticipating any corresponding phenomenal patterns. So, the covariance theory is inapplicable to imaginative phenomenal intentionality.

I am assuming that the red spherish content of my bizarre imaginative episode is literally identical to the red spherish content of my orderly perceptual episode. But perhaps a reductionist will deny this. Perhaps she will insist that unless you expect your imaginative episode to evolve according to the anticipatory conditionals characteristic of red rubber balls, you aren't imagining a red rubber ball. Instead, you are representing something else that that has some features in common with a red rubber balls. And states that represent those things – red *shrubber* balls, let's call them – are characterized by a distinctive set of anticipatory conditionals. Your imaginative state *does* include the expectation of systematic covariance among phenomenally-presented features. It just doesn't include the expectation of the systematic covariance characteristic of red-rubber-ball-representing states.¹⁸

I don't find this move plausible. Why can't I imagine a red rubber ball behaving radically differently from how I expect it would in real life? More to the point:

since represented items can behave in any old way in the imagination, it is hard to see how content is to be individuated in the imagination. Are red shrubber balls things that behave in *this* particular chaotic way, but not in a slightly different, but still chaotic, way? Suppose I imagine a ball behaving normally and then suddenly behaving chaotically. Did the kind of thing represented by my imaginative state change, or was I imagining a non-rubber, non-shrubber ball the whole time?

Even if these questions can be answered satisfactorily, there is a simpler type of content exhibited by the imagination that is not plausibly associated with any anticipatory conditionals, viz., two-dimensional items such as a red square. When I imagine a red square, I have no expectations whatsoever about how it will behave, because I need not imagine it as embedded in a dynamic manifold of any kind. So it cannot be that implicit anticipatory conditionals associated with my imaginative state make it the case that squarehood is part of its intentional content: there are no such conditionals associated with it.

Premise 3: The phenomenal intentionality of imaginative states is to be explained in the same fashion as that of perceptual states. Even if I am correct that the covariance theory cannot explain imaginative phenomenal intentionality, this would pose no threat to the covariance theory if the covariance theory only applies within the domain of perceptual experience. But if it applies within the domain of perceptual experience, it ought to apply to the imaginative as well, because we should expect that phenomenal intentionality in imaginative states is of a piece, metaphysically, with phenomenal intentionality in perceptual states. To see why, try the following: focus visually on some nearby object. Then look away while holding its appearance in mind. Finally, imagine what it would look like for it to be oriented slightly differently. You have just experienced a smooth transition from visual perception through visual memory to visual imagination. While a change has occurred in the *phenomenal mode* within which you consciously represent the object, no change has occurred regarding the *phenomenal vehicle*, viz., that familiar visual manner of conscious representation. Correspondingly, it is hard to make sense of the claim that, somewhere during the transition, a sharp break occurred regarding the way that your phenomenology comes to represent properties of the object to which you initially visually attended. Note, further, that perceptual-intentional contents and imaginative-intentional contents can overlap and be mistaken for one another, per the so-called 'Perky effect.'¹⁹ Thus it would be ad hoc to posit a deep metaphysical difference between perceptual phenomenal intentionality and imaginative phenomenal intentionality, in the absence of strong and independent theoretical reasons to do so.

In sum: if the covariance theory were true of phenomenal intentionality within the perceptual domain, it would be able to explain imaginative phenomenal intentionality. But it cannot, so it fails in both domains.

4. Second objection to reductionism: explanatory gappiness

My second objection is that, even within the perceptual domain, there is an explanatory gap between the awareness of phenomenal patterns, on the one hand, and phenomenal intentionality, on the other. It is not transparent how the former could constitute the latter. There are at least two explanatory questions that the covariance theory doesn't answer, one regarding phenomenal representation and the other regarding phenomenal content: (1) Why would awareness of phenomenal patterns constitute the presentation of intentional contents, rather than not? (2) Supposing they do so constitute, what makes it the case that particular objects and properties are presented, rather than others?

As regards (1), compare: the pattern of brushstrokes in a painting constitutes the painted image. And the painted image depicts something (e.g. a pond with lilies). But of course the patterns do not, in themselves, depict anything. If I recognize the painting as a picture of a pond, that is because I, the viewer, visually interpret it thus, or I learn that the artist intended it thus. So it is hard to make sense of the claim that phenomenal patterns as such constitute intentional states.

Now, Masrour seems to anticipate this objection, at least as the objection applies to phenomenal objectivity (that is, to the alleged constitution by phenomenal patterns of *mind-independent* representational contents). He writes:

[P]henomenal properties enter phenomenal concepts as constitutive components. Accordingly, whether the phenomenal concepts of phenomenal objectivity and schematic dynamical unity are a priori independent depends on whether the phenomenal properties that they pick out coincide in us, and our observations support that they do. Thus the two concepts are not a priori independent when we consider them as phenomenal concepts.²⁰

Masrour's reply is that if (a) a phenomenal state's presenting intentional contents *just is* (or is constituted by) the awareness, on the part of the subject, of the phenomenal patterns exhibited by that state, and (b) phenomenal concepts are individuated by the very phenomenal state-types they're about (a familiar if controversial claim from the literature on phenomenal concepts), then there's something muddle-headed about asking *how* the awareness of phenomenal patterns, on the one hand, constitutes phenomenal objectivity, on the other. These are co-referring concepts, after all!

I don't think this response will work, because even if the two concepts in question are *partly* individuated by their referents, they are not *wholly* individuated by their referents. That is, my concept of the awareness of phenomenal patterns has more content to it than 'the phenomenal state-type: _____', where the blank is filled in with an inwardly-ostended phenomenal token. Asking how the awareness of phenomenal patterns could constitute phenomenal objectivity is not merely asking how tokens of *that* phenomenal state-type could constitute tokens of *that* phenomenal state-type. It is more like asking how aggregates of H₂O molecules could constitute water, where we have some grasp of the

concepts of H₂O molecules and of water. That's a good question, and it hasn't been answered until the constitutive relationship has been rendered intelligible. Reductionists owe us an intelligible account of the constitutive relationship between non-intentional components of phenomenal states and the intentional features they allegedly constitute.

Even more daunting is the second explanatory question: what makes it the case that particular objects and properties are presented, rather than others? Suppose that two super-cognizers were able to keep a running list of all the somatosensory properties instantiated by a human subject over a sufficiently long time-interval. Suppose further that the super-cognizers were informed (a) that some of the patterns of instantiation fix intentional content; and (b) which patterns in particular do so. And then suppose they were asked to *interpret* the current phenomenal state of the subject. How divergent a pair of interpretations might the two super-cognizers come up with? I say: *massively* divergent. While the first super-cognizer might treat the subject as representing persisting objects in her environment, the second might treat the subject as representing gappy objects scattered throughout space and time, or temperature changes on the surface of the sun. My worry, in other words, is that phenomenal patterns could be interpreted as representing spatiotemporally local states of affairs, but they could also be interpreted as representing infinitely many spatiotemporally distal states of affairs.

A defender of the covariance theory might protest: it is precisely *perceptual* phenomenal intentionality that phenomenal patterns constitute, so non-localist interpretations can be ruled out. But I don't think we should be satisfied by this response. What makes it the case that phenomenal patterns constitute *perceptual* phenomenal intentionality rather than memories of perceptual experiences, or imaginative experiences that honor the same regularities we encounter in perception? But even if we require that our super-cognizers treat the relevant patterns as constituting perceptual states, they could still interpret these perceptual states in a range of ways. For example, it is a datum that perceptual experiences often present *enduring* objects. And there is nothing in the phenomenal patterns themselves that would demand being interpreted in terms of enduring objects, rather than in terms of very short-lived objects – or, for that matter, in terms of property-instances inhering in indiscrete gunk.

So the covariance theory, at least as developed by Farkas and Masrour, fails to transparently explain the constitution of phenomenal intentionality by non-intrinsically intentional phenomenology. It doesn't directly follow, of course, that every attempt to reduce phenomenal intentionality to a set of somatosensory ingredients is doomed to fail. Nevertheless, it is hard to see what an alternative version of reductionism would look like. To provide a constitutive explanation of some macro-level feature, a would-be reductionist has two resources to draw upon, broadly speaking: (a) features of the components of the macro-level entity that has the feature in question, and (b) relations among these components.

(Recall: to provide a reductive explanation of the climbability of a ladder, one has to show how the features of the parts, in tandem with the relations among those parts, give rise to climbability.) The covariance theory takes just this shape. Certainly no theory that makes use of *fewer* resources is going to succeed. But it's also hard to see what could be added. We have already considered a variety of types of relations among sensorimotor properties, including covariance-patterns at a time, covariance-patterns over time, subsets thereof, the subjects' awareness thereof, the subjects' anticipations of further iterations thereof. None of these refinements have helped. I have not *proven* that there are no alternative reductionist strategies waiting to be developed. Still, a provisional verdict against reductionism, and in favor of primitivism, seems warranted.

4. Primitivism about phenomenal intentionality

I propose the following theory of phenomenal intentionality. Among the phenomenal properties are *phenomenal-intentional properties*, or '*P-I* properties'. *P-I* properties are irreducible to any other phenomenal properties; they comprise a *sui generis* type of phenomenal property, whose most abstract determinable is *being phenomenally-intentionally directed some-content-wise*.²¹ The nature of a *P-I* property wholly consists in the presentation to the subject of some particular intentional object. An example is the *P-I* property whose intentional object is *causation*. When a subject instantiates this property, she is thereby consciously presented with *causation*. Similar things can be said about other intentional contents that are necessitated by the phenomenal character of a conscious mental state.

P-I properties are what phenomenal duplicates share that explains intentional sameness (per phenomenal duplication arguments). For example, what explains how my envatted twin and I can share phenomenally- and thus intentionally-identical perceptual states as of an object dragging another object is that we both instantiate the *P-I* property whose object is causation (in connection with a number of other *P-I* and somatosensory properties). Likewise, *P-I* properties are what intentional non-duplicates *fail* to share that explains the phenomenal difference between them (per intentional contrast arguments). For example, what explains why I hear the sentence 'the cause of the war was vengeance' as a claim about the war's *intended purpose* and my envatted twin hears the sentence as a claim about the war's *causal origin* is that his conscious mental state includes the *P-I* property whose intentional object is causation, whereas mine includes the *P-I* property whose intentional object is goal-directedness.

The arrival of primitivism on the scene brings with it good news and bad news. The good news is that primitivism is not subject to the objections discussed above in connection with reductionism. First, the very same *P-I* properties can be included in conscious mental states of different types. What makes it the case that a perceptual state is about one thing's causing another can

make it the case that an *imaginative* state is about one thing's causing another. This is good news because it explains how perceptual phenomenology and imaginative phenomenology can generate intentional contents in the same way, despite the fact that we anticipate predictable changes within perceptual phenomenology but often do not anticipate any sort of predictability within imaginative phenomenology.

Second, primitivism does not require the positing of an enigmatic constitution-relation that is said to hold between phenomenal patterns, on the one hand, and intentional properties, on the other. Rather, some phenomenal properties *by their very nature* confer intentional content on states that include them. Now, questions remain regarding the type of explanatory relation at play in phenomenal-intentional states. Why does the inclusion of a *P-I* property in a phenomenal state entail that that state has certain intentional properties? One answer to this question is that the explanatory relation in question is *identity*. This is the view, for example, of Angela Mendelovici (2010) and David Bourget (2010). On this view, there can be no phenomenal properties that lack intentional contents, and there can be no intentional properties – at least, of the 'original' or 'underived' sort – that lack a phenomenology.

I reject identity-primitivism for two reasons. First, identity-primitivism is a form of (non-reductive) representationalism, according to which all phenomenal properties are intrinsically representational.²² But I deny that somatosensory properties intrinsically represent. Second, it is perfectly sensible to speak of the instantiation of intentional properties in the absence of phenomenal properties. For example, metaphysicians debate whether certain abstracta, such as propositions or Fregean senses, intrinsically represent what they're about. But surely there is nothing it is phenomenally like *for a proposition* to represent some state affairs. So, I propose that the relevant explanatory relation is not identity but rather *realization*. Intentional properties are realized by *P-I* properties, though there are other ways for them to be realized. (I'll say a bit more about this in a moment.)

That's the good news. Alas, the arrival of primitivism also heralds bad news. Notoriously, primitivisms give us everything we want at no cost. Primitivisms are, or can be, inquiry-stoppers, amounting to the 'reification of a puzzle rather than its solution.'²³ If, in the famous example, we 'explain' opium's soporific properties in terms of a primitive *virtus dormitiva*, our inquiry actually *loses* ground, because we forestall a proper investigation of the chemical properties of opium. Human representational capacities are, presumably, considerably richer and more complicated than the causal capacities of opium, worthy of sustained scientific and philosophical analysis.

There are really two worries here. The first is a quite general worry about primitivisms of all sorts. I quite agree that primitivist 'explanations' can be vacuous and anti-scientific. Now, it cannot be the case that primitivism is *always* illegitimate: some features of the world are fundamental, after all, and we do

well to acknowledge them as such when we come across them. But lest we mistake a derivative feature of the world for a fundamental one, the burden of proof really is on the primitivist. In particular, she needs to show, first, that reductionist alternatives inevitably fail, or at any rate that we have very good reason to expect that they will fail. (Such was the task I took up in the previous two sections.) Second, she needs to show that her proposed ontological expansions do real explanatory work – that they render intelligible otherwise puzzling phenomena within a domain and/or that they reveal illuminating connections between domains. I think that just such a comprehensive, positive case can be made for *P-I* properties. There is no space here to *mount* such a case. But I will try to sketch, at the end of this section, what such a case could look like.

The second worry is not about whether primitivisms in general are theoretically vicious but about whether phenomenal-intentional primitivism in particular is appropriate. Specifically, the worry is that primitivism is too clean and simple – it papers over the limitations and idiosyncrasies of our intentional capacities. For one thing, it might be thought that primitivism requires a crazy profusion of primitive phenomenal properties. We can consciously think about any individual or category we can demonstrate or assign a name to; yet it is implausible that there is a primitive, proprietary phenomenal property corresponding to every such intentional content. Call this ‘The Abundance Problem.’²⁴ For another, if the contents of every intentional state of ours were available to us as the intentional object of a *P-I* property, we ought to be able to fully ‘grasp’ those contents, i.e. to know what those properties essentially consist in. But the difficulty (some would say *failure*) of conceptual analysis shows that this is not the case: our grasp of our thought-contents is quite poor. Call this ‘The Analysis Problem.’²⁵

I will attempt to address both worries – the general worry and the particular worry in its two forms – by explaining the *crucial but limited role* that I take *P-I* properties to play in a theory of intentionality. I will use as a jumping-off point the semantic theory developed by Bertrand Russell in ‘Knowledge by Acquaintance and Knowledge by Description.’ For Russell, propositional attitudes are polyadic relations between a subject and the various constituents of a proposition. If, for example, the relation *judging* holds between a subject, Ann, love, and Bill (in that order), then the subject judges that Ann loves Bill.²⁶ Russell also holds that we can understand propositions only to the extent that we are ‘acquainted’ with their constituents: ‘Every proposition which we can understand must be composed wholly of constituents with which we are acquainted.’²⁷ Russellian acquaintance is a relation that holds between a subject and an object of one of two types: a mental particular (he calls them ‘sense-data’) or a universal.

So far, the theory is unable to account for myriad judgments we make about objects with which we are unacquainted (in Russell’s sense). Unless Ann and Bill are mental particulars, it would seem that the subject cannot make judgments about them, since she cannot be acquainted with them. And this isn’t just an

issue for particulars; Russell acknowledges further that there are universals with which we are unacquainted, but about which we can make judgments.

Famously, Russell's solution is to introduce descriptions into the picture. In order to make a judgment about an object with which we are unacquainted (either a particular or a universal), a subject needs to be able to construct a definite description of that object out of objects with which the subject *is* acquainted. To use Russell's example, if an associate of Bismarck wished to make a judgment about him, he might construct a description out of certain sense-data 'which he connected (rightly, we will suppose) with Bismarck's body.' We who have never met Bismarck will have to employ objects further removed from him: 'the description in our minds will probably be some more or less vague mass of historical knowledge – far more, in most cases, than is required to identify him.'²⁸

I reject a number of elements of Russell's picture, but I endorse Russell's core semantic intuition: all intentional contents are either directly grasped by a subject or else indirectly picked out via something like a description constructed from directly grasped elements. With *P-I* properties in the picture, we have a way to endorse Russellian semantics without endorsing Russell's metaphysics. *P-I* properties present an intentional content to the subject. They can thus explain how it is that a subject 'directly grasps' an intentional content. And if they are able to semantically combine with one another, they can form phenomenal complexes expressive of descriptions. Call this neo-Russellian view 'phenomenal descriptivism.'

We are now in a position to see how intentional properties can be *realized* by *P-I* properties, but can also be realized in other ways. Take some intentional property, *being about F*. One way for it to be realized is for a subject to instantiate a *P-I* property whose intentional object is *F*; the subject's phenomenal state thereby realizes the intentional property *being about F*. A second way for this intentional property to be realized is for the subject to instantiate a structured complex of *P-I* properties that expresses a description that picks out *F*. If the description so expressed contains an indexical or demonstrative (e.g. 'that thing there with such-and-such features'), the description will conscript into semantic service whatever bit of extra-mental reality answers to the demonstrative or indexical. In such cases, the realizer for the intentional property *being about F* will include bits of extra-mental reality. There are thus as many potential ways for *being about F* to be realized as there are ways for the subject's mental state to be about *F*. It may be the case that most of the intentional properties I instantiate are realized in part by extra-mental goings-on, and some may be realized by an enormous amount of extra-mental goings on. (Example: 'Let's fill our glasses with *the favorite beverage of the kings and queens of the country Julia wanted to visit last year.*') The only *P-I* properties a subject instantiates are those whose intentional objects provide the basic, descriptive content needed to conscript the world into semantic service.

Thus the primitivism I advocate regarding *P-I* properties is a *sparse* primitivism. And sparse primitivism is able to avoid the Abundance Problem and the Analysis Problem mentioned in section four. As regards the Abundance Problem: There is no proprietary *P-I* property for every possible thought-constituent; rather, *P-I* properties form a sparse set, combinations of which serve to pick out countless kinds and particulars.²⁹ As regards the Analysis Problem: it is no surprise that conceptual analysis is elusive, so long as intentional contents can be fixed by indexical- or demonstrative-involving modes of presentation, and hence fixed in part by the extra-mental goings on conscripted by those modes of presentation. In sum: primitivism about *P-I* properties does not commit us to a simplistic picture of conscious content. Rather, *P-I* properties provide the basic building blocks of an enormous, complicated phenomenal-intentional edifice.

In the remainder of this section I return to the first, more general concern about primitivism. Primitivisms earn their keep, I grant, only when they provide deep and broad theoretical payoff. And there are indeed a number of theoretical roles that *P-I* properties can play outside of the context in which I have invoked them (viz., explaining how phenomenology necessitates intentionality). First, *P-I* properties can help solve puzzles in epistemology. Consider the problem of *semantic self-knowledge*. It is a datum that are usually in a position to *know what we mean* when we think thoughts or linguistically express ourselves. Such an epistemic achievement is interesting in its own right; in addition, it plausibly explains a host of other epistemic abilities we have, such as knowing when two words are synonymous, knowing whether we have adequately captured what we were trying to say, knowing whether a perceptual experience satisfies a linguistic description ('Aha – I seem to be looking at the round, purple tree Alice was telling me about yesterday'), and so forth. Now, once *P-I* properties are in the picture, we can readily understand semantic self-knowledge as a species of *introspective* knowledge: we know what our thoughts mean the same way that we know what our sensations feel like, viz., by introspectively attending to them. But without *P-I* properties in the picture, it is much harder to understand what semantic self-knowledge could be. Some philosophers have attempted to reduce semantic self-knowledge to one or another epistemic ability, such as our ability to know when two thoughts have the same content, or our ability to trace out the logical consequences of our thoughts.³⁰ The trouble is that semantic self-knowledge plausibly explains a *bunch* of abilities, not just one. And a theory of semantic self-knowledge consisting in a *list* of such abilities isn't satisfying: lists of phenomena don't make for good theories, at least not when it's possible to account for all the items on the list by appeal to a single explanatory posit (in our case, introspectable *P-I* properties).³¹

Another problem in epistemology which *P-I* properties can help solve is the problem of *rational insight*. It seems that we are able to know the truth of some propositions simply by reflecting on them; such reflection affords the subject not only a conviction of their truth but genuine insight into what makes

them true.³² An example discussed by Chudnoff (2011) is the proposition *every diameter of a circle determines a line of symmetry for it*. 'When the proposition ... appears to you to be true, it does so in a way that seems grounded in your awareness of what circles, diameters, and lines of symmetry *are* – that is, in a felt presence to mind of the properties of being a circle, being a diameter, and being a line of symmetry.'³³ The trouble is figuring out what such 'presence to mind' could come to. It is typically taken to require that we be able to cognitively apprehend causally inert Platonic abstracta. But because this is widely believed to be impossible, many philosophers reject rational insight altogether.

P-I properties provide a better way of making sense of 'presence to mind'-talk, and hence a way to take rational insight seriously. Suppose I instantiate *P-I* properties whose intentional contents are circularity, being a diameter, and being a line of symmetry. I am thus presented with these contents within the manifold of my own mind; I can compare them, combine them, manipulate them, etc. (assuming that the requisite cognitive capacities are intact). Now, I bear no *direct* cognitive relation to the universals *circularity*, *being a diameter*, and *being a line of symmetry*. But that's not a problem. Neither do I bear a direct cognitive relation to the objects of perception. In the case of perception, so long as my perceptual phenomenal-intentional states stand in the right dependence-relation to the objects of perception (in this case, non-deviant causal dependence), I bear an indirect cognitive relation – the *perceiving* relation – to those objects. Likewise: so long as the *P-I* property whose intentional object is *circularity* stands in the right dependence-relation to the universal *circularity*, I stand in an indirect cognitive relation – let's call it the 'apprehension' relation – to that universal. Thus, when I compare, combine, and manipulate *P-I* properties, I can have rational insight into the relations among the universals that serve as the intentional objects of those properties.^{34,35}

But there's more. Suppose that phenomenal descriptivism could serve not only as a theory of conscious intentionality but as a theory of *all* intentionality. Here's the idea: I have been suggesting that the intentional contents of a *conscious* mental state are fixed by (a) that state's *P-I* properties + phenomenal indexicals/demonstratives, in tandem with (b) whatever bits of extra-mental reality are conscripted into semantic service by (a). Phenomenal descriptions thus serve as *modes of presentation* of all the intentional contents of conscious mental states. But now suppose that a subject's non-conscious mental states (such as standing beliefs) just are dispositions to instantiate *conscious* mental states. In that case, *P-I* properties would supply the basic semantic mechanism that fixes the content of all of subject's intentional states.³⁶ I cannot of course develop and defend this proposal here. But if it proves viable, the case for *P-I* properties would be strong indeed.

6. Conclusion

Let us take stock. The strategy (as developed by Farkas and Masrour) of reducing phenomenal intentionality to phenomenal patterns has ontological parsimony on its side. Nevertheless, it is inapplicable to imaginative phenomenal intentionality, and we have just as strong of reasons to believe in imaginative phenomenal intentionality as to believe in perceptual intentionality. Further, it rests on an obscure explanatory proposal: that of the *constitution* of intentionality by phenomenal patterns. Attempts to respond to these objections on behalf of Farkas and Masrour have not succeeded, suggesting that reductionism itself is the problem (rather than their proposals' peculiarities). Phenomenal primitivism, at least of the *sparse* sort, avoids these problems, while holding out the promise of explanatory rewards in other domains. Primitivism, despite its requirement that we expand our inventory of basic phenomenal properties – is therefore to be preferred.

Notes

1. The phrase 'phenomenal intentionality' comes from Horgan and Tienson (2002). My usage differs from theirs in two ways. First, they introduced the phrase as the name of a *thesis*, rather than of a *phenomenon*. The thesis in question is, 'There is a kind of intentionality, pervasive in human mental life, that is constitutively determined by phenomenology alone.' Second, the phenomenon I am calling 'phenomenal intentionality' is weaker than the one mentioned in their thesis: I am characterizing the phenomenon in terms of a necessitation-relation, rather than in terms of a determination-relation. Necessitation is weaker than determination, because necessitation, unlike determination, (a) can be symmetrical, and (b) does not entail that one relatum explains the other.
2. Ibid. As Horgan & Tienson stress, perceptual states may have *wide* accuracy conditions that are not introspectively manifest, but their *narrow* accuracy conditions – which partly determine wide accuracy conditions – *are* introspectively manifest.
3. Siegel (2010), p. 72.
4. Loar (2003). More precisely, such experiences share intentional contents as of lemony-looking things; this is consistent with saying that subjects need to be in causal contact with lemon-tokens in order to represent the natural kind *lemon*.
5. Siewert (1998), p. 321.
6. Philosophers who reject cognitive-phenomenal intentionality include Lormand (1996), Prinz (2011), Tye and Wright (2011), and William Robinson (2005, 2011). Philosophers who reject even perceptual-phenomenal intentionality include Lycan (2008) and Carruthers and Veillet (2011).
7. In a recent article, Walter Ott (2016) purports to fill an explanatory lacuna in the phenomenal-intentionality literature by appealing to resemblance. If I am reading him correctly, the explanandum of interest to him is phenomenal reference. For my own part, I do not expect that there is such a thing as phenomenal reference.
8. An alternative form of reductionism proceeds in the opposite direction: it attempts to reduce phenomenal properties to intentional properties. According to so-called 'reductive representationalist' theories of phenomenal properties, all

phenomenal properties are constituted by intentional properties (in tandem, it is sometimes suggested, with their functional role). Reductive representationalists will obviously not be motivated by the appeal to parsimony that I just made, since they deny that there are any non-intentional phenomenal properties. But I am going to set reductive representationalism aside in what follows; my interest is in what we might call 'phenomenology-first' theories of phenomenal intentionality.

9. Farkas (2013), p. 109.
10. *Ibid.*, p. 100. Note that this passage could be read as an account of phenomenal representation or phenomenal objectivity, rather than of phenomenal content. Elsewhere she claims that our 'image of the world is constructed from basic experiential features' (note 6, p. 114). Since our image of the world involves a rich palate of properties, she must be thinking that phenomenal structure necessitates phenomenal content, in addition to the more general features of phenomenal representation and phenomenal objectivity.
11. It is not uncontroversial that such properties are ever instantiated in our experience. Representationalists deny that they are (see footnote 8 above). But reductionists about phenomenal properties need them, since they need non-intrinsically intentional phenomenal properties to form the reduction-base.
12. 'The feeling as you hold the book, the small noise the paper makes as you turn a page, all these features form a structure that, in itself, and also *together with past experiences*, clusters the experienced features around the well-known objects presented in the manifest image of our world' (*Ibid.*, 108, emphasis mine).
13. Masrour (2013) does claim (a) that there is a kind of intentionality that is constituted by phenomenality, and (b) that he thinks that the conditions under which a state exhibits phenomenal objectivity are precisely those conditions under which a state exhibits intentionality that is constituted by phenomenality. So, if he means by 'intentionality that is constituted by phenomenality' roughly what I mean by 'phenomenal content,' – an interpretation that strikes me as natural – then Masrour believes that his reductive theory of phenomenal objectivity is *eo ipso* a reductive theory of phenomenal content.
14. Masrour (2013), p. 132.
15. *Ibid.*, 122.
16. She calls it the 'sensory' domain. If there is any *cognitive* phenomenal intentionality – and she is sympathetic to the idea – it is 'basic and not constructed.' (2013, 113).
17. This is not to say that imaginative contents in general float free of one's external situation. I am able to imagine my brother playing golf (say) only because my representation of my brother is causally connected to him. My life-long envatted twin, when he undergoes an imaginative episode phenomenally identical to mine, does not consciously represent my brother. Nevertheless, there are *some* imagined contents – generic red sphericity, for example – that are necessarily shared by phenomenally identical twins.
18. Thanks to a referee for suggesting this response on behalf of the reductionist.
19. Perky (1910).
20. Masrour (2013), 133.
21. Compare Kriegel (2011), ch. 3, who seems to have a similar view in mind.
22. This view differs from *reductive* representationalism (see footnote 8) in leaving unreduced the phenomenology of phenomenal intentionality. To say that every phenomenal property is identical to an intentional property is not, by itself, to attempt a reduction of either property-type.
23. The line comes from Kirk Ludwig (2012); his target is a primitivist move in the philosophy of language.

24. Wilson (2003) and Tye and Wright (2011) raise versions of this concern.
25. See Pitt (2011).
26. Russell gives his example using the letters 'A' and 'B'.
27. Russell (1910), 117.
28. *Ibid.*, 114–115.
29. Which contents are in the set? Research in cognitive linguistics and developmental psychology suggest at least the following, for us humans: object, agent, cause, motion, egocentric location, and some set of shape- and size-properties. Is that all? Well, I suspect they can take us pretty far, if they are used in *metaphorical* descriptions as well as literal ones. (See Camp 2006). Perhaps the set of primitive contents needs to be considerably more abundant. Even so, it will contain a tiny fraction of all of the contents we can entertain.
30. Cf. Millikan (1993), Burge (1988) and Peacocke (1996).
31. I read Pitt (2004) as making the same suggestion, viz., that phenomenally-grounded intentionality is needed to explain certain of our introspective discriminatory capacities. See also Bourget (2017) who argues, much more extensively than I can here, that a phenomenological theory of mentally 'grasping' contents is explanatorily preferable to an inferentialist theory (according to which one's grasping of *p* is reducible to one's ability to infer the consequences of *p*).
32. Writes Laurence Bonjour (*forthcoming*, 3.): 'Rational insight penetrates beyond or beneath the proposition as a whole to reveal just how and why various component properties and relations and the overall structure combine in such a way that the resulting proposition has to be true.'
33. Chudnoff (2011), 637.
34. Now, the dependence-relation that holds between universals and the intentional objects of corresponding *P-I* properties is not causation, of course, but rather a non-causal determination-relation. Which determination-relation in particular? The answer will depend on one's metaphysics of universals. A believer in immanent universals could say that the relation is *partial constitution*. This seems to be Bonjour's view in his (1998): triangularity-thoughts, he tentatively suggests, involve the instantiation of a 'complex universal having the universal triangularity as one of its components' (184). For a different account of the relevant determination-relation, see Chudnoff (2013).
35. 'Are *P-I* properties really less mysterious than direct cognitive relations to Platonic abstracta?' Well, I think so (though it's not clear quite how to measure and compare degrees of metaphysical mystery). For one thing, it's not clear that it's even possible for a direct cognitive relation to hold between a temporal mind and an atemporal universal. For another, there are ways of explaining which *P-I* properties get instantiated by a subject at a time, but I'm not sure there is a way of explaining which universals are apprehended at a time. (How is one's Platonic apprehension-faculty supposed to get *aimed* at the right universal?) Thanks to a referee for pressing me here.
36. Searle (1992) seems to hold this view, or something very close to it. See also Gertler (2007), Strawson (2008), ch. 11, and Chalmers (2012), 466–467. I develop a positive proposal in my (2016).

Disclosure statement

No potential conflict of interest was reported by the author.

Notes on contributor

Philip Woodward is an assistant professor of Philosophy and a member of the Neuroscience Program committee at Valparaiso University. He received his PhD from Indiana University in 2015.

ORCID

Philip Woodward  <http://orcid.org/0000-0002-6768-9477>

References

- Bonjour, Laurence. *forthcoming*. In *Defense of Rational Insight*. <http://faculty.washington.edu/bonjour/Unpublished%20articles/BonJour,%20In%20Defense%20of%20Rational%20Insight.pdf>.
- Bonjour, Laurence. 1998. In *Defense of Pure Reason*. New York, NY: Cambridge UP.
- Bourget, David. 2010. "Consciousness is Underived Intentionality." *Noûs* 44 (1): 32–58.
- Bourget, David. 2017. "The Role of Consciousness in Grasping and Understanding." *Philosophy and Phenomenological Research* 95 (2): 285–318.
- Burge, Tyler. 1988. "Individualism and Self-Knowledge." *Journal of Philosophy* 85 (11): 649–663.
- Camp, Elisabeth. 2006. "Metaphor in the Mind: The Cognition of Metaphor." *Philosophy Compass* 1 (2): 154–170.
- Carruthers, Peter, and Benedicte Veillet. 2011. "The Case against Cognitive Phenomenology." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 35–56. New York, NY: Oxford UP.
- Chalmers, David. 2012. *Constructing the World*. Oxford: Oxford UP.
- Chudnoff, Elijah. 2011. "What Intuitions are Like." *Philosophy and Phenomenological Research* 82 (3): 625–654.
- Chudnoff, Elijah. 2013. "Awareness of Abstract Objects." *Noûs* 47 (4): 706–726.
- Farkas, Katalin. 2013. "Constructing a World for the Senses." In *Phenomenal Intentionality*, edited by Uriah Kriegel, 99–114. New York, NY: Oxford UP.
- Gertler, Brie. 2007. "Overextending the Mind?." In *Arguing about the Mind*, edited by Brie Gertler and L. Lawrence, 192–206. Abingdon: Routledge.
- Horgan, Terence, and John Tienson. 2002. "The Intentionality of Phenomenology and the Phenomenology of Intentionality." In *Philosophy of Mind: Classical and Contemporary Readings*, edited by David Chalmers, 520–533. New York, NY: Oxford UP.
- Kriegel, Uriah. 2011. *The Sources of Intentionality*. New York, NY: Oxford UP.
- Loar, Brian. 2003. "Phenomenal Intentionality as the Basis of Mental Content." In *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, edited by Martin Hahn and Bjørn Ramberg, 229–258. Cambridge: MIT Press.
- Lormand, Eric. 1996. "Nonphenomenal Consciousness." *Noûs* 30 (2): 242–261.
- Ludwig, Kirk. 2012. "What Role Should Propositions Have in the Theory of Meaning? Review Essay: Scott Soames. What is Meaning?" *Philosophia* 40 (4): 885–901.
- Lycan, William. 2008. "Phenomenal Intentionalities." *American Philosophical Quarterly* 45 (3): 233–252.
- Mendelovici, Angela. 2010. *Mental Representation and Closely Conflated Topics*. Doctoral Dissertation.

- Millikan, Ruth Garrett. 1993. "Knowing What I'm Thinking of." *Proceedings of the Aristotelian Society* 67: 91–124.
- Masrour, Farid. 2013. "Phenomenal Objectivity and Phenomenal Intentionality: In Defense of a Kantian Account." In *Phenomenal Intentionality*, edited by Uriah Kriegel, 116–134. New York, NY: Oxford UP.
- Peacocke, Christopher. 1996. "Entitlement, Self-Knowledge, and Conceptual Redeployment." *Proceedings of the Aristotelian Society* 96: 117–158.
- Ott, Walter. 2016. "Phenomenal Intentionality and the Problem of Representation." *Journal of the American Philosophical Association* 2: 131–145.
- Perky, C. W. 1910. "An Experimental Study of Imagination." *The American Journal of Psychology* 21: 422–452.
- Pitt, David. 2004. "The Phenomenology of Cognition or What is It like to Think That P?" *Philosophy and Phenomenological Research* 69 (1): 1–36.
- Pitt, David. 2011. "Introspection, Phenomenality, and the Availability of Intentional Content." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 141–173. New York, NY: Oxford UP.
- Prinz, Jesse. 2011. "The Sensory Basis of Cognitive Phenomenology." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 174–196. New York, NY: Oxford UP.
- Robinson, William. 2005. "Thoughts without Distinctive Non-Imagistic Phenomenology." *Philosophy and Phenomenological Research* 70 (3): 534–562.
- Robinson, William. 2011. "A Frugal View of Cognitive Phenomenology." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 197–214. New York, NY: Oxford UP.
- Russell, Bertrand. 1910. "Knowledge by Acquaintance and Knowledge by Description." *Proceedings of the Aristotelian Society* 11 (5): 108–128.
- Searle, John. 1992. *The Rediscovery of the Mind*. Cambridge: MIT Press.
- Siegel, Susanna. 2010. *The Contents of Visual Experience*. New York, NY: Oxford UP.
- Siewert, Charles. 1998. *The Significance of Consciousness*. Princeton UP: Princeton.
- Strawson, Galen. 2008. *Real Materialism and Other Essays*. Oxford: Oxford UP.
- Tye, Michael, and Briggs Wright. 2011. "Is There a Phenomenology of Thought?" In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 326–344. New York, NY: Oxford UP.
- Wilson, Robert A. 2003. "Intentionality and Phenomenology." *Pacific Philosophical Quarterly* 84 (4): 413–431.
- Woodward, Philip. 2016. "Conscious Intentionality in Perception, Imagination, and Cognition." *Phenomenology and Mind* 10: 140–155.