

A Structuralist Theory of Phenomenal Intentionality

Ben White

Forthcoming in *Inquiry*

<https://doi.org/10.1080/0020174X.2023.2242398>

Abstract: This paper argues for a theory of phenomenal intentionality (herein referred to as ‘Structuralism’), according to which perceptual experiences only possess intentional content when their phenomenal components are appropriately related to one another. This paper responds to the objections (i) that Structuralism cannot explain why some experiences have content while others do not, or (ii) why contentful experiences have the specific contents that they have. Against (i), I argue that to possess content, an experience must present itself as an experience of something distinct from itself, and that only experiences whose components are suitably structured satisfy this requirement. Against (ii), I argue that Structuralists can provide a deeper explanation of perceptual content than other theories of phenomenal intentionality by showing how the content-determining relations among our experiences depend on the selection pressures under which our perceptual systems developed.

Keywords: phenomenal intentionality; mental content; phenomenal consciousness; perception; perceptual content.

1. Introduction

Theories of phenomenal intentionality (PIT) hold that there is a form of mental content that is determined by phenomenal consciousness (Searle 1992; Horgan and Tienson 2002; Loar 2003; Strawson 2010; Kriegel 2013). Advocates of PIT differ, however, in how they construe this claim. Identity PITists claim that every experience (i.e. phenomenally conscious mental state) has an intentional content that is *identical* to its phenomenal character (Bourget 2010; Mendelovici 2018). Primitivist PITists hold that not all experiences have content, but those that do have the specific contents that they have as a matter of brute fact (Woodward 2019). Still others, whom I’ll call Structuralist PITists (or Structuralists), hold that experiences only have content when their phenomenal components are appropriately related to one another (Farkas 2013; Masrour 2013).

This paper defends Structuralism against two objections that have been raised by advocates of the other forms of PIT just mentioned: (i) that Structuralism cannot explain why some experiences have content while others do not, and (ii) that it cannot explain why contentful experiences have the specific contents that they have.

In response to (i), Sections 2 and 3 expand on versions of Structuralism proposed by Farkas (2013) and Masrour (2013) by clarifying the nature of the relations that must obtain among an experience's components for it to have intentional content. In response to (ii), Sections 4 and 5 argue that Structuralists can provide a deeper explanation than other forms of PIT of why contentful experiences have the specific contents that they do by showing how the content-determining relations among our experiences depend on the selection pressures under which our perceptual systems developed.

Two quick preliminaries: First, while some PITists aim to account for the content of all intentional mental states, I'll be focusing solely on perceptual content. Henceforth, the term 'experience' will therefore be used to refer strictly to *perceptual* experiences. I'll also use the terms 'experience' and 'phenomenal state' interchangeably throughout.

Second, as the primary aim of this paper is to defend Structuralism against rival forms of PIT, I won't be defending PIT against theories that treat mental content as determined solely by non-phenomenal factors. Most prominent among such theories are causal and teleological theories of content (henceforth 'externalist theories'), which hold, roughly, that a mental state represents F if and only if it causally covaries with F in a certain way, or has the biological function of indicating F (Fodor 1987; Dretske 1988; Millikan 1989; Neander 2017). In making such causal or biological relations the sole determinants of mental content, externalist theories have difficulty explaining our ability to form distinct, determinate representations of necessarily coincident entities that have

the same effects on our sense organs (Hall 1990; Adams and Aizawa 1994; Gates 1996; Strawson 2010; Horgan and Graham 2012). PIT avoids this problem by giving the requisite content-determining function to phenomenal consciousness. According to PIT, the aforementioned ability thus follows from the fact that the same proximal stimuli can produce experiences with different phenomenal characters, which endow those experiences with different contents. Much more would need to be said to make these points convincing, but that is a task for another paper. Here I will simply assume, on the strength of these considerations, that *if* (*contra* Quine (1960) and Dennett (1987)) our intentional mental states have determinate content, this is because such content depends on phenomenal consciousness. The aim of this paper is to show that of the various forms of PIT, Structuralism provides the best way of developing this idea.

2. Why Structuralism?

In arguing for Identity PIT, Mendelovici (2018, 95, 101) raises several objections to Structuralism, the most central of which is that Structuralists and other PITists who deny that *all* experiences have intentional content must explain why some experiences have content while others do not. As she puts it: ‘Absent specific reasons to think that there is such a significant division between phenomenal states, the view that all phenomenal states give rise to intentionality seems more likely.’ While Primitivist PITists can justly be charged with dodging this question (since they hold that what content, if any, an experience has is a brute fact that admits of no further explanation), Structuralists can provide a more satisfying response by arguing that in order to have content, an experience must present itself as an experience of some object or property (henceforth simply ‘thing’) distinct from itself, and that experiences that are too simple or chaotic fail to satisfy this constraint, as they provide no basis for the experiencing subject to distinguish independently

existing entities presented in their experience from features of their experience itself. Consequently, only experiences that are both sufficiently complex and organized can have content, as only in such cases do the phenomenal components of the experience present themselves not (or not merely) as features of the experience, but rather as things that exist independently of the experience in question.

Considered in isolation, simple experiences like color sensations, the feeling of heat, or the taste of bitterness seem incapable of supporting any such distinction unless combined with other experiences in an appropriate spatio-temporal structure (Farkas 2013, 100). Taken by itself, apart from any other experiences occurring before, after, or simultaneously with it, a simple experience of red thus doesn't seem to represent anything distinct from that experience as being one way instead of another. Likewise for a chaotic swarm of sensations that is highly complex, but too disorganized to be taken as indicating anything about how things are independently of itself.

One might suppose that a uniform experience of red across one's visual field just represents the property 'redness.'¹ I take it, though, that for an experience to *represent* a property as opposed to merely *instantiating* it, the experience must present the represented property as existing independently of itself. In the case of properties like redness (or heat, or bitterness), this seems to require that the experience present the property as instantiated or inhering in some independently existing object or region that is distinct from the experience in question; at the very least, it must represent something like 'redness at such-and-such independently existing location.' Perhaps a completely homogenous experience of red might then be said to represent something like 'red

¹ Mendelovici (2018, 42-5) suggests that simple experiences like this represent what Chalmers (2006, 66) calls 'Edenic' properties, i.e. 'simple intrinsic qualities of objects in the world [...] whose nature [(i.e. the qualities')] we seem to grasp fully in perceptual experience.' Importantly, these Edenic properties are supposed to be perceptually represented as objective properties of objects. If the above line of thought is correct, however, then we can't represent the Edenic properties that Mendelovici treats such simple experiences as representing without presupposing a capacity for objective representation that such experiences are too simple to possess by themselves.

everywhere.’ But whence do we acquire the representation of objective space expressed in the term ‘everywhere’? Taken in isolation from the more complex experiences in which they figure, such simple experiences seem incapable of supporting any such representation. A being whose experience consisted solely of an unchanging uniform experience of red thus could not acquire from it any representation of objective space. Likewise for an experience consisting solely of a chaotic swarm of different colors. Though highly complex, such an experience lacks the regularity and structure needed for the components of the experience to present themselves as features of something that exists independently of the experience itself. As a result, the experience fails to be *about* anything distinct from itself, and thus lacks content. A being whose experience consisted solely of such a chaotic swarm of colors thus again couldn’t acquire from it any representation of objective space, or of any independently existing entities whatsoever.

Why is it, then, that experiences that *are* appropriately structured thereby represent things independent of themselves? As PITists, Structuralists must hold that this is ultimately due to the phenomenal character of such experiences, which differs in an all-important way from that of experiences that are too simple or chaotic to have content. In short, there is something about what it’s like to have an experience whose components are suitably organized that makes such experiences present themselves as experiences of independently existing entities in a way that unstructured experiences do not. More informatively, Structuralists might add that suitably structured experiences have this distinctive character *because* the regular patterns in such experiences make it seem as though they aren’t simply random or self-generated, but are instead being produced by certain stable, enduring objects and properties that exist independently of the experience, and which are responsible for (and thus explain) its non-random character. As a result, the experience presents itself as an experience *of* those objects and properties that the experience

seems to be produced by or under the control of. Note that this is not a point about what's *actually* causing the experience; even a highly structured, contentful experience might be completely hallucinatory. Nor is it an epistemic point about the beliefs one acquires or the judgments that one makes or is disposed to make in response to such an experience. Though having an experience whose components are appropriately structured might give one a *pro tanto* reason to take it at face value and judge that it is being produced by something independent of itself, one needn't make or be disposed to make any such judgment for the experience to have content. An appropriately structured experience can thus have the phenomenal character of having been produced by something distinct from itself and thereby have content even if the experiencing subject is firmly convinced that they're hallucinating. The point is instead a purely phenomenological one about the phenomenal character of the experience itself, viz. that due to the regular patterns among their components, such experiences *appear* to us as though they are being produced by things that are independent of themselves, which they thereby present to us as the objects of our experience, regardless of whether we judge or are disposed to judge that they have in fact been so produced. This is why Structuralists treat an experience's structure as playing such a crucial role in determining what content, if any, it has, and why they differ from Primitivists in thinking that it's not entirely brute or inexplicable why suitably structured experiences have content while overly simple or chaotic ones do not. For it is only in virtue of some regular pattern among its components that an experience presents itself as produced by, and thus as informing us of (or presenting us with) things independent of itself.

3. Which relations confer content?

The foregoing considerations motivate the Structuralist view that the reason why only some experiences have content is that an experience must be both sufficiently complex and organized to present itself as an experience of something distinct from itself. But what precisely is the nature of the relations among an experience's components that determine what content, if any, it has? Obviously, these must be whatever relations are missing in experiences that are either too simple or too chaotic to have content. But can we say anything more specific about them?

Building on the prior proposals of Farkas (2013) and Masrour (2013) (who argue on similar grounds that our experiences acquire content by virtue of certain systematic relations they bear to one another, with special importance placed on how they change in response to our own perceived bodily motions), I suggest that the relations that confer content on our experiences can be divided into two sorts. First are what I'll call basic structural relations. These are spatio-temporal relations (e.g. an experience of yellow occurring immediately to the left of another concurrent color experience in one's visual field, or a tactile experience occurring before, after, or simultaneously with an auditory experience) and relations of qualitative similarity (e.g. an olfactory experience being qualitatively indistinguishable from some prior olfactory experience, or wholly dissimilar to a tactile experience). These can be conceived as organizing the phenomenal components of any complex experience into a 'mosaic' in which each component or 'tile' occupies a particular spatio-temporal location and stands in a network of similarity relations to every other tile.

Second are what I'll call grouping relations. These include relations of segmentation and grouping produced by factors like Gestalt principles, expectations, and mechanisms of perceptual constancy. Such factors can be roughly conceived as inducing a function from a set of experiences structured by the basic structural relations to that same set of experiences further organized into

distinct perceptual objects. For example, given a particular ‘mosaic’ of experiences ordered by the basic structural relations:

- the Gestalt principles of proximity and similarity might group certain adjacent, qualitatively similar color experiences in the mosaic together into a single perceptual unit;
- the Gestalt principles of closure and good continuation might further group two separate but colinear line segments of similar color together as segments of the same, partially occluded edge, and the principle of common fate might group similarly moving components of an experience together as parts of the same object;
- certain expectations may cause the components of an ambiguous figure to be organized in one way instead of another (e.g. by segmenting two lines into a ‘1’ and a ‘3’, instead of grouping them together as a ‘B’);
- lightness constancy might group two adjacent but dissimilar color experiences together as parts of the same uniformly colored surface (the change in color being attributed to the presence of a cast shadow), or size constancy might present a shrinking color region as an object that’s receding into the distance instead of diminishing in actual size.

Though I’ve suggested that one can conceive of such grouping relations as being generated by the application of Gestalt principles, expectations, constancy mechanisms, and other such factors to a mosaic of phenomenal states that has already been partially structured by the basic structural relations, this is a bit misleading, for the grouping relations among an experience’s components can alter the basic structural relations among them, e.g. by increasing the qualitative similarity and/or spatio-temporal proximity of components that are grouped together as part of the same perceptual object, or, conversely, by increasing the dissimilarity and/or spatio-temporal distance between components that are segmented into two distinct objects. The Kanizsa triangle illustrates

the latter effect. Due to the way that the color experiences on either side of the illusory edges of the triangle are grouped (in this case, via the Gestalt principles of good continuation and closure), people experience the inside ‘edge’ of the triangle as brighter than (and thus as qualitatively dissimilar to) the background that it is perceived as occluding, despite the fact that there is no corresponding physical difference in the figure itself (Palmer 1999, 292). Hence, although I may sometimes speak as if the grouping relations are superimposed on a set of experiences that has already been organized by the basic structural relations, the basic structural and grouping relations are in fact produced in tandem, with each constraining and influencing the other. I also stress (i) that the grouping relations among an experience’s components are part of the experience’s phenomenal character, and (ii) that while the grouping relations are constrained by the basic structural relations, they are not fully determined by them. Experiences consisting of the same components standing in the same basic structural relations to one another can thus have different phenomenal characters because their components are grouped differently. Gestalt switches provide an illustration of this point.

Adopting Farkas’ (2013) terminology, I’ll describe a set of experiences as being ‘externalized’ when the basic structural and grouping relations among them are such as to give them the character of being produced by something distinct from themselves, which they thereby represent. What kinds of relations are the appropriate, externalization-enabling ones? In our case, they seem to include (a) high degrees of qualitative similarity among spatio-temporally contiguous experiences, (b) smooth contours among similarity-based groupings (so that sharp discontinuities in qualitative similarity form spatio-temporally continuous outlines, as in the line formed by a color boundary, or the motion of such a boundary in a connected path across one’s visual field), and (c) cross-modal agreement in similarity-based groupings (so that the relations among experiences in one

sensory modality correspond in an orderly way to those among experiences in other sensory modalities). Evidence of the importance of (a)-(c) can be drawn from studies of object perception in infants, cross-modal interactions and illusions, and principles of perceptual organization in vision and audition, which suggest that stimuli that produce experiences whose components are thusly related are thereby more likely to be perceived as persisting objects, or single environmental sources of sound (Spelke 1990; Bregman 1993; Palmer 2003, 179-87; O'Callaghan 2012; Casati 2015).

The foregoing list is not meant to provide a set of necessary and sufficient conditions for the externalization of experiences. Examples of externalized experiences that exhibit some but not all of the aforementioned relations thus come readily to mind. When listening to a musical recording, one's auditory experience may be externalized by virtue of the relations of qualitative similarity and continuity among temporally adjacent experiences so as to represent a certain stream of sound as a melody produced by a single instrument, even if there are no concurrent changes in other sensory modalities that correspond to the changes in the melody, in which the case the externalized experience might be said to exhibit relations (a) and (b), but not (c). That said, (a)-(c) should serve to indicate the kinds of relations that are reliably present in contentful experiences, e.g. the correlations between the tactile and visual experiences one has when looking at an object that one is rotating in one's hands, or between the visual and auditory experiences one has while looking at a speaker's face.

A few quick clarificatory points: As I'm using the term, 'externalization' is not a further act of judgment that an experiencing subject must make to endow their experiences with content. Rather, to say that an experience is externalized is just to say that, due to the relations among its components, it presents itself as an experience of something distinct from itself. As stressed in the

previous section, this is a purely phenomenal feature of the experience itself, not an additional epistemic state or disposition that the experience produces in the experiencing subject. Also, while Mendelovici and Woodward describe Structuralism as the view that certain relations among experiences ‘give rise to’ intentional content, I think Structuralists are better off holding that an experience’s content (assuming it has any) is *constituted by* (but not identical to²) its phenomenal character, which itself consists of the phenomenal character of its components and the basic structural and grouping relations among them (and, in the case of contentful experiences, the appearance of having been produced by something distinct from the experience itself). In addition to avoiding difficult questions about the nature of the ‘giving rise to’ relation, this alternative has three advantages.

First, it enables Structuralists to allow (as Identity PITists cannot) that experiences with different phenomenal characters can nevertheless possess the same content, by arguing that just as it is possible for a statue to be constituted by a different lump of clay than the one that now constitutes it, so too the content possessed by a given experience could also be constituted by an experience consisting of different components and/or different relations among them.

Second, it enables Structuralists to explain (as Identity PITists may have a hard time doing) why we ascribe different properties to an experience when speaking of its content than we do in relation to its phenomenal character. Just as certain properties (e.g. the property of being Romanesque, or of being well made) might be properly ascribed to a statue but not to the clay that constitutes it (and *vice versa*)³, so too Structuralists might argue that certain properties of an experience’s content (e.g. having certain satisfaction conditions, or being accurate or inaccurate)

² While some (e.g. Lewis (1971), Noonan (1993), and Biro (2018)) have argued that constitution is a form of identity, the reasons for thinking otherwise are to my mind sufficiently plausible to make this a viable position for Structuralists to adopt (Johnston 1992; Baker 1997; Fine 2003).

³ See Fine (2003).

are not aptly ascribed to its phenomenal character, and conversely, that certain properties of an experience's phenomenal character (e.g. feeling a certain way) are not properly attributable to its content.

Third, since constitution is plausibly a metaphysically necessary relation (in that if certain entities constitute y , then it is metaphysically impossible for y 's constituents to bear those same relations to one another without constituting y), this enables Structuralists to accommodate the intuition (which some PITists see as a major motivation for their view⁴) that there is a form of content that any phenomenally indiscernible experiences necessarily share in common. One could of course achieve these same results by postulating a large class of brute, metaphysically necessary 'phenomenal-intentional' laws dictating that experiences consisting of such-and-such components related in such-and-such a way thereby 'give rise to' a certain content, but unless an explanation can be given for why these laws have this modal status, this seems rather *ad hoc*.

Putting all this together, Structuralists can respond to Mendelovici's objection by arguing that to possess intentional content, an experience must be about something distinct from itself, and that experiences only become externalized in this way when their components bear the right kind of basic structural and grouping relations to one another, where the appropriate relations (in our case) include (a)-(c).

4. The biological basis of phenomenal structure

In arguing for his preferred, Primitivist version of PIT, Woodward (2019, 615) raises another objection to Structuralism in the form of two questions that he claims it fails to answer: '(1) Why would awareness of phenomenal patterns constitute the presentation of intentional

⁴ See Horgan and Tienson (2002, 524-7) and Horgan and Graham (2012, 335-7).

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?’ The preceding discussion provides an answer to (1): Certain patterns among phenomenal states (e.g. those involving relations (a)-(c)) endow experiences with content because experiences whose components are organized in such a way thereby present themselves as experiences of things that exist independently of our experience of them. I further suggested in Section 2 that the reason *why* suitably structured experiences present themselves in this way is that the regular patterns among their components make such experiences seem as though they’re not random or self-generated, but are instead being produced by some stable, enduring entities distinct from themselves, which such experiences thereby inform us of. This, however, leaves (2) unanswered. For even granting that certain relations among phenomenal states endow our experiences with content, one might wonder why experiences whose components are organized in a particular way thereby represent one thing instead of another.

Here it first bears noting that Primitivism is not the only form of PIT that treats the connection between phenomenal characters and contents as in some sense brute. Indeed, *any* PITist account of content must bottom out at some point in the claim that experiences with *this* phenomenal character have *this* content as a matter of brute fact. This is the cost of holding that what an experience represents is determined ultimately by what it’s like to have it. The difference between Structuralism and Primitivism is thus *not* that Structuralists think that some further explanation can be given for why experiences with a given phenomenal character have the specific content that they have, but rather that Structuralists don’t think it’s entirely inexplicable why some experiences have content while others do not. For the Structuralist, there is a striking (and *explanatory*) difference in the phenomenal characters of contentful and contentless experiences, viz. that contentful experiences are sufficiently complex and organized that they do not appear to

be random or self-generated, but instead present themselves as having been produced by things distinct from themselves.

This doesn't imply, however, that some further explanation can be given for why a certain phenomenal structure correlates with the specific type of content that it does (i.e. an explanation that goes *beyond* the claim that there's something it's like to have an experience whose components are organized in this way, which phenomenal character by nature confers that content on any thusly organized experience). Indeed, Structuralists *can't* do this without abandoning PIT! For in order to avoid simply adverting once again to the nature of the experience's phenomenal character, such an explanation would have to appeal to something *beyond* that character to explain why it gives the experience the specific content it has. But this is effectively to say that the experience's content is ultimately determined by something *besides* its phenomenal character. For if the character of the experience can't determine its content on its own without this additional, non-phenomenal factor to explain *why* that character is correlated with this content, then the additional, non-phenomenal factor would seem to be what's *really* determining the content of the experience. And that is inconsistent with PIT.

As PITists, Structuralists are thus entitled to answer Woodward's second question by saying that the reason why an experience whose components are organized in a certain way thereby has the particular content it has is due ultimately to the fact that there's something it's like to have an experience whose components are thusly organized, and this phenomenal character is by nature such as to endow any such experience with that type of content. Indeed, as PITists, they have no choice but to say this, and to insist that they say *more* than this is to demand of them something that no PITist can do.

Nevertheless, there is also a sense in which Structuralists *can* say more than other PITists about why our experiences have the kind of contents that they do. For while Structuralists must, like all PITists, treat the explanation of an experience's content as terminating ultimately in certain brute correlations between phenomenal characters and contents, for Structuralists, any such explanation will be incomplete without some account of how our experiences came to be organized in the way they typically are. Here Structuralists can do more than just point to the phenomenal character of our experiences. By appealing to the selection pressures under which our perceptual systems developed, they can provide a deeper explanation of why experiences instantiating certain basic structural relations thereby tend to be grouped in certain ways as opposed to others, and thus have the general type of content that they do.

This, however, involves a change in the *kind* of explanation being offered. In particular, this additional explanation doesn't replace or modify the foregoing account of what ultimately *constitutes* or otherwise determines an experience's content. On this issue, Structuralists must, like all PITists, insist that an experience's content is ultimately determined by its phenomenal character alone. In *this* sense, the causes of the experience and the evolutionary/learning history of the experiencing subject have no bearing on what the experience represents. That said, it would, I think, be a mistake for Structuralists (or anyone for that matter) to claim that such etiological factors are *completely* irrelevant to experiential content. For even if they aren't part of what *constitutes* an experience's content, they can still help to explain why it has the phenomenal character (and hence the content) that it has, by explaining why the components of the experience are grouped in the way that they are.

As just mentioned, though, this will be an explanation of a very different sort; i.e. not a *constitutive* explanation, but rather a *historical* one, which appeals to the evolutionary/ecological

history of the experiencing subject to show why, in perceivers with this type of history, experiences consisting of such-and-such phenomenal components standing in such-and-such basic structural relations to one another are apt to be grouped in this way. By itself, the constitutive explanation of experiential content outlined above doesn't address this question. *Given* that an experience consists of such-and-such components standing in such-and-such relations to one another, the constitutive explanation tells us that it thereby has a distinctive phenomenal character whose nature is such that it constitutes whatever content the experience has. This, however, does nothing to explain *why* the experience's components are grouped as they are. It is *this* question that the historical explanation provides an answer to.

As we'll see below, there are limits to how far the historical explanation can go in answering this question. In particular, while it can explain why, e.g., certain components of an experience are grouped in such a way as to represent something edge-like, it *cannot* explain why those components are so grouped as to represent an edge instead of an edge stage, undetached edge parts, or any other gavagai-ish edge variants. The relevant contrast class for the historical explanation is thus at a more general, abstract level than the fully determinate content of the experience. Claims about the historical explanation's ability to explain why the components of an experience are grouped in a certain way should be understood with this point in mind.

Given the impact that Structuralists treat the grouping of an experience's components as having on its phenomenal character and thus its content, in explaining why an experience's components are grouped in a certain way, there is a sense in which the historical explanation thereby shows the content of our experiences to be partly determined by non-phenomenal factors concerning our evolutionary/learning history. While this may seem incompatible with PIT, it is not. For viewed as an explanation of experiential *content*, the historical explanation differs from

the constitutive explanation in being both *partial* and *indirect*. It is *partial* because it explains only that aspect of an experience's content that is constituted by its components being grouped in one way as opposed to another (and even here, the explanation is incomplete in the sense described in the previous paragraph). It is *indirect* in that the evolutionary/ecological factors cited in the historical explanation don't constitute or directly determine an experience's content, but instead affect it *only* via the impact that they have on its phenomenal character (viz. by explaining why its components are grouped in a certain way). As PITists, Structuralists will again insist that whatever etiological factors are involved in producing an experience's phenomenal character, it is this character (however produced) that alone constitutes or otherwise directly determines the experience's content. With these points in mind, Structuralists can thus allow that the non-phenomenal, evolutionary/ecological factors cited in the historical explanation help to explain (albeit in a partial and indirect way) experiential content without thereby abandoning PIT.

Having clarified the difference between these two forms of explanation and how they fit together in the Structuralist account of experiential content that I'm developing, I'll now present the historical component of the account in greater detail.

The primary function of our perceptual systems is to help us successfully navigate our environment by extracting information about our surroundings from sensory input. This task is complicated by the fact that sensory input is ambiguous, as there are countless different environmental scenarios that could have produced the proximal stimuli that we receive at any given moment. While theorists differ over the extent and importance of this ambiguity, the standard view is that our perceptual systems disambiguate sensory input by relying on various 'assumptions' that enable them to rule out certain potential causes of the stimuli we're currently

receiving, thereby winnowing down the potential ways that the environment might be.⁵ These assumptions capitalize on contingent regularities in the environment and have been incorporated into our perceptual systems through learning or natural selection. In the case of vision, examples include the assumption that objects are typically rigid, opaque, and locally continuous in depth, that illumination typically comes from above, and that abrupt changes in the texture or reflectance of a surface are relatively rare (Palmer 1999, 218, 228, 244, 489). Such assumptions underlie the Gestalt principles, perceptual constancies, and depth cues by which our visual experiences are organized.

While it may seem at odds with the kind of personal-level, phenomenological account of intentionality that PIT aims to provide, I think Structuralists can use these points regarding the sub-personal processes that produce our experiences to augment their constitutive account of experiential content with a historical account of the sort described above, thereby providing a more comprehensive answer to Woodward's second question.⁶ More specifically, Structuralists can hold that when the components of a complex experience E are so related that E thereby has a particular content, this is partly owing to the fact (a) that the grouping relations among E's components correspond to certain biologically significant relations among the stimuli that E was produced in response to, and (b) that, due to the selection pressures under which they developed, our perceptual systems have incorporated assumptions that lead them to convert such stimuli into experiences that are grouped in such a way as to have the kind of content that E has.⁷ In short, the grouping

⁵ A noteworthy departure from this orthodoxy is the ecological approach to perception championed by Gibson (1979).

⁶ Like many 'why' questions, Woodward's second question can be read as ambiguous between a historical question ('How did it come about that experiences whose components are arranged in such-and-such a manner have the contents they do?') and an ahistorical one ('Regardless of the historical processes that produced this result, what is it about this kind of phenomenal structure that gives such experiences the contents they now have?'). Taken together, the constitutive and historical explanations address both of these readings.

⁷ Note that the fact that experiences with E's structure have such content is itself presupposed, not explained, by the historical explanation. This fact is accounted for only by the constitutive explanation, which the historical explanation thus depends upon.

relations that help determine the phenomenal character and thus the content of our experiences depend on the assumptions that our perceptual systems have adapted to use in processing sensory input.

Suppose, e.g., that light reflected from a surface and its occluded background is projected onto my retina, resulting in a visual experience that represents an edge at a certain location. My visual representation of the edge consists of an externalized boundary in my visual field defined by a discontinuity in the color experiences on either side of it. The spatial and similarity relations among the experiences that constitute this boundary correspond to certain spatial and similarity relations among the stimuli that produced them; experiences occupying adjacent regions in my visual field are produced by light falling on adjacent regions of the retina, and experiences that are qualitatively similar in color are produced by light with similar spectral characteristics.⁸

As a result, the edge-representing boundary in my experience corresponds to a discontinuity in the spectral characteristics of the light that produced the experience. Because such discontinuities in the visual stimulus are regularly correlated with biologically significant facts concerning the size, shape, and location of edges, our visual system has adapted to convert such stimuli into experiences containing color boundaries that are grouped in such a way as to represent edges.⁹ The perceived size, shape, and location of any such edge thus depends on the size, shape,

⁸ This isn't to say that similarity relations among visual experiences are perfectly isomorphic to similarity relations among stimuli; metamerism, the 'filling in' of the blind spot where the optic nerve leaves the retina, and the influence of memory color, chromatic adaptation, and simultaneous chromatic contrast on perceived color are just a few examples of ways in which similar color experiences may be produced by stimuli with different spectral characteristics (Delk and Fillenbaum 1965; Palmer 1999, 34, 102-3, 106-8, 616-7; Hansen et al. 2006). To keep things manageably simple, I'll bypass such complications here. My proposal requires merely that such isomorphism be sufficiently advantageous to favor the development of perceptual systems that regularly ensure that biologically significant relations among stimuli correspond to content-determining grouping relations among the experiences they produce in response to such stimuli.

⁹ See footnote 7. Here again, the fact that such groupings constitute representations of edges itself lies beyond the reach of the historical explanation. It is instead addressed only by the constitutive explanation, which the former explanation thus depends upon.

and location of the color boundary in the visual experience that represents it, which features of the color boundary depend in turn on the spectral discontinuities in the stimulus.¹⁰

If our environment or visual apparatus had differed in such a way that spectral discontinuities in the visual stimulus *didn't* regularly correlate with edges, then our visual system wouldn't have adapted to respond to these discontinuities in the way that it has. It would instead rely on different 'assumptions' in processing visual input, and the experiences that our visual system produced would consequently be structured by *different grouping relations*. An experience whose components *would* be grouped and externalized as representing an edge in our current condition might thus fail to be externalized in a counterfactual scenario wherein our visual system developed under different selection pressures.

Returning, then, to the proposal that this example is meant to illustrate, the basic idea is as follows: Experiences whose components stand in certain basic structural relations are typically grouped in the way that they are (and thus have the phenomenal character and content that they do) in part because the groupings in question correspond to certain relations in the stimulus that the experience was produced in response to. These relations in the stimulus are (or were at some point) regularly correlated with some biologically significant feature of the environment, and because of this, our perceptual systems have adapted to convert such stimuli into experiences whose components are so grouped as to represent a feature of that type. Content-determining groupings among the components of an experience are thus dependent on the relations that our

¹⁰ In addition to spectral discontinuities of the sort described above, there are of course many other cues (e.g. T-junctions, accretion and deletion, the concavity/convexity of contours, binocular disparity and convergence) that our visual system uses to determine whether such discontinuities are produced by surface edges, and if so, which side of the edge is the occluding figure and which the occluded ground (Palmer 1999, 205-25, 229, 237, 281-4, 290). In the absence of these additional cues, the visual system is less likely to convert a spectral discontinuity into an experience in which the color boundary corresponding to that discontinuity is grouped and externalized as an edge (e.g. when the boundary is produced by a cast shadow). This is because our visual system developed in an environment in which such discontinuities are more regularly correlated with edges when these additional cues are also present.

perceptual systems have adapted to pick up on in the stimulus that the experience was produced in response to. The resulting account of perceptual content can be divided into four stages:

- (i) Due to contingent regularities in our environment, biologically significant features of our surroundings are regularly correlated with certain relations among proximal stimuli.
- (ii) Since it is biologically advantageous for one's perceptual systems to pick up on such correlations, our perceptual systems have adapted to process sensory input on the basis of the assumption that the environment exhibits the regularities that these correlations depend on.
- (iii) Because our perceptual systems have incorporated these assumptions, the basic structural and grouping relations among our experiences correspond to certain relations among the proximal stimuli that our perceptual systems produce them in response to (which in turn correspond to certain biologically significant features of the distal environment).
- (iv) Because an experience has the phenomenal character that it does (consisting of the phenomenal characters of its components, the basic structural and grouping relations among them, and, in the case of suitably organized experiences, the character of being produced by something independent of the experience itself), it has the content that it has.

The explanatory relation between *explanans* and *explanandum* differs across these four stages. In stages (i)-(iii), which make up the historical explanation described above, the *explanans* nomically necessitates the *explanandum*, although the relevant laws differ in each case. In (i), the relevant laws are those that enable us to determine an individual's proximal stimuli from their current surroundings. In (ii), the relevant laws are those that undergird the mechanisms of learning and natural selection. In (iii), the relevant laws are the neurophysiological laws that determine the patterns of neural activity that our perceptual systems produce in response to various stimuli,

perhaps together with psycho-physical laws that determine the kind of experience (if any) that a given neural state realizes.

The constitutive explanation kicks in at stage (iv). While one *could* treat the explanatory relation in (iv) as one of nomic necessitation by postulating a class of *sui generis* 'phenomenal-intentional laws' whereby experiences possessing certain phenomenal characters thereby 'give rise to' certain intentional contents, for reasons mentioned above, Structuralists are I think better off holding that an experience's content is instead constituted by (but not identical to) its phenomenal character. Either way, the role that (iv) assigns to an experience's phenomenal character in explaining its content distinguishes the overall account as a form of PIT. The emphasis on the role that the relations among an experience's components play in determining its content at this stage further distinguishes it as a form of Structuralism.

As the fact that a given structure produces certain effects explains, in other cases, why that structure was preserved under some selection process, there is likewise a sense in which stages (i)-(iii) in the foregoing explanation depend upon (iv). Just as the fact that organs with the structure and makeup of a human heart are capable of pumping blood explains why such organs were preserved under natural selection (given the biological advantages we gain by having an organ with this capability), so too the fact (cited at stage (iv)) that experiences whose components are organized in a certain way thereby have a certain content explains why our perceptual systems developed to respond to certain stimuli by producing experiences whose components are thusly organized (given the biological advantages we gain by having experiences whose contents are generally accurate, together with the fact that the content possessed by thusly organized experiences is generally accurate in such stimulus conditions). However, the fact that thusly organized experiences *do* possess such content itself cannot be explained by the historical

explanation comprised of (i)-(iii), but is instead presupposed by it. As PITists, Structuralists must again hold that this fact receives its explanation only at stage (iv), with the claim that the phenomenal character of experiences whose components are so organized is by nature such as to give them the content in question. The historical explanation thus depends ultimately upon the constitutive.

5. Why not externalism?

Those sympathetic to externalist theories of content of the sort mentioned in Section 1 may wonder why stage (iv) is necessary. Once the basic structural and grouping relations among experiences have been shown to covary with environmental features as described in stages (i)-(iii), isn't that sufficient to show why experiences whose components are related in such-and-such a manner have the contents that they do *without* invoking their phenomenal character? Perhaps the only role that phenomenal consciousness plays in all this is to give subjects an awareness of the contents of their perceptual states, which have *already* been fixed by the physiological and evolutionary/ecological factors mentioned in (i)-(iii).

While the aim of this paper is not to defend Structuralism or PIT in general against externalism, a few words are nevertheless in order to explain why stage (iv) *is* necessary from the standpoint of the Structuralist position I'm developing here. First, it bears re-emphasizing that on this view, the basic structural and grouping relations are themselves *part* of phenomenal character of the experiences whose content they help constitute. Consequently, the link that (i)-(iii) establishes between these relations and distal features of the environment already extends to the phenomenal level. Phenomenal consciousness has already entered the picture by stage (iii).

Second, and more importantly, from the Structuralist's perspective, were it not for the fact that the basic structural and grouping relations *are* phenomenal features of our experiences (and that the phenomenal states they relate are likewise phenomenally conscious), there would be no fact of the matter as to what our experiences represent. For given any environmental feature F that a certain complex of basic structural and grouping relations among experiences happens to covary with, those relations will covary just as regularly with many other environmental features that are distinct from but necessarily coincident with F. Insofar, e.g., as the relations among the components of the experience described in the previous section are regularly correlated with the presence of edges, they will *ipso facto* be equally correlated with edge stages, undetached edge parts, and countless other gavagai-ish edge-variants. Any biological benefits derived from having experiences whose components are thusly organized in the presence of edges will consequently be equally served by having such experiences in the presence of edge stages, undetached edge parts, etc. From a strictly externalist perspective, there are thus no grounds for treating such experiences as representing any one of these features instead of the others. This is the moral of the indeterminacy objections to externalism mentioned in Section 1. While the assumptions built into our perceptual systems enable them to disambiguate sensory input to some degree, if we view those assumptions as determined solely by environmental regularities that have had a causal impact on ourselves or our phylogenetic ancestors, then there will always be many different ways of representing a stimulus that they'll fail to discriminate between.

Hence the need for stage (iv). From this perspective, only if we take into account the impact that the grouping relations involved in the aforementioned experience have on its phenomenal character can we explain why such experiences represent edges as opposed to edge stages, undetached edge parts, or the disjunction of such things. For this difference is one that only shows

up at the phenomenal level, viz. in the fact that experiences whose components are grouped in such a way as to represent an edge have a *different phenomenal character* than experiences whose components are grouped differently so as to represent an edge stage, or undetached edge parts. While it could've been equally biologically advantageous for our perceptual systems to structure our experiences in any of these alternative ways, it so happens that the experiences they produce in the presence of edges and their gavagai-ish kin are in fact grouped in such a way as to represent *edges* (where this grouping is again a phenomenal feature of the experiences in question, which explains *why* they have the determinate contents that they have).

Though I've been focusing on the impact that the grouping relations among an experience's components have on its phenomenal character and content, the phenomenal components that instantiate these relations of course contribute to the experience's phenomenal character and content as well. Experiences whose components stand in the same relations to one another may thus differ in content insofar as the components themselves have different phenomenal characters. Given two visual experiences whose components stand in the same basic structural relations to one another and are so grouped as to represent a surface of a certain shape, size, and location, one experience may represent the surface's color as red, while the other represents its color differently, because the phenomenal components grouped and externalized as representing the surface in the one case are phenomenally red, and in the other case have some other phenomenal character.

As with the basic structural and grouping relations, here too the mechanisms of learning and natural selection may be indifferent to the specific phenomenal character of the simple experiences or 'qualia' that our experiences happen to be composed of. Given a set of alien color qualia that bear the same similarity relations to one another and are correlated with the same spectral stimuli as the components of our phenomenal color space, from an evolutionary/ecological

perspective, it'd make no difference which type of phenomenal states our visual experiences are composed of, since both would be equally useful in enabling us to distinguish objects in terms of their spectral properties.¹¹ Here too, there thus may've been a number of paths our perceptual systems could have gone down, each of which would've been equally good from an evolutionary/ecological standpoint, and it so happens that we've taken the path that leads to our representing the colors of objects via our current color qualia, instead of via some other, structurally isomorphic set of simple experiences.

For PITists, this is again not just a point about the phenomenal character of our color perceptions; it affects their content as well. If our visual experiences were composed of the alien color qualia just described, we would perceptually represent objects as having different colors than those we now perceive them as having. This difference in content could not, however, be captured by externalist theories, as these alternative color experiences would *ex hypothesi* stand in the same similarity relations to one another, covary with the same stimuli, and hence be just as biologically useful in distinguishing surfaces and detecting spectral differences as our actual color experiences. Here too PITists will thus see a need to treat the phenomenal character of the components of our experiences as playing an additional role in determining their content. This is why *both* the phenomenal characters of an experience's components and the basic structural and grouping relations among them are cited in stage (iv) as constituting its content.

Putting stages (i)-(iv) together, we have an explanatory chain running from contingent environmental regularities to experiential content that provides a more comprehensive answer to Woodward's second question than that provided by the constitutive explanation in stage (iv) alone. Like any form of PIT, the proposed Structuralist account admittedly retains an element of bruteness

¹¹ Barring, of course, any further differences in the physiological conditions and associated costs/benefits of producing these different experiences.

in the connection between an experience's phenomenal character and its content, as indicated by the fact that (iv) is in a sense all that's needed to explain why an experience has the content that it does. Unlike other forms of PIT, though, the proposed account at least helps us see why certain structural features of our experiences have the kind of impact on their content that they do, why certain experiences are externalized while others are not, and why the components of any experience are apt to be grouped in the way that they are, so that the resulting experience represents certain types of features instead of others. In conjunction with stage (iv), stages (i)-(iii) thus explain why the components of the experience described in the previous section are grouped in such a way as to represent an edge or some gavgai-ish variant thereof instead of some completely different feature (e.g. weight, pitch, hardness, or price). Like all forms of PIT, though, certain questions about an experience's content (e.g. why experiences structured like ours represent *edges* instead of edge stages, or why appropriately structured experiences involving sensations of redness represent objects as *red* as opposed to some other color) can on this view be answered only by saying that that's the kind of content that experiences with that kind of phenomenal character by nature have. Hence, again, the need for stage (iv).

That stages (i)-(iii) in the foregoing account nevertheless *do* add something to the explanation of experiential content is indicated by the counterfactual dependence of (iv)'s *explanandum* on these three preceding stages. Thus, as noted above, if the environment in which our perceptual systems developed had been sufficiently different, they would have adapted to use different assumptions in processing sensory input. Experiences produced in response to a given stimulus would consequently be grouped differently, and would thereby have a different phenomenal character and hence content than they do.¹²

¹² The kittens in Blakemore and Cooper's (1970) experiment provide an illustration of this point. This isn't to say that *only* beings with a biological history can have contentful experiences, though. I see no reason why a robot or

Objection: ‘If fully determinate content only emerges at the personal, phenomenological level, where does this leave the various “assumptions” that are cited in stages (i)-(iii) as explaining why our experiences are organized in the way they are and thus possess the content that they do? If these assumptions are operative only at the unconscious, sub-personal level of perceptual processing, then it seems that they can have no determinate content (i.e. there can be no determinate fact of the matter as to what exactly our perceptual systems are “assuming”). To put it another way, either (i) our perceptual systems *really do* possess the determinate contents that we ascribe to them in treating them as implementing certain assumptions about the environment, or (ii) any such ascriptions should be understood non-literally, as part of a useful but fictitious way of talking about the sub-personal causes of our experiences. If (ii), then how can the “assumptions” built into our perceptual systems explain anything about the content of our experiences, given that they don’t actually possess the content that such explanations treat them as having? If (i), then given that sub-personal states of our perceptual systems can possess determinate content despite being phenomenally unconscious, why must we appeal to the phenomenal character of our experiences to explain why they have the contents they do? Why aren’t the assumptions built into our perceptual systems sufficient to determine experiential content on their own?’

This objection can be met by distinguishing between two different but equally legitimate ways of viewing our perceptual systems and the assumptions they implement. We can view our perceptual systems either (a) solely in light of the selection pressures that shaped them, and treat them as assuming the existence of regularities in the environment that they have adapted to pick up on. Or (b), we can view our perceptual systems in light of the experiences they produce, and

Swampman couldn’t have experiences with the same contents that ours have. The explanation of why their experiences are organized in the way that they are will differ, however, from the corresponding explanation for biological creatures like us.

treat them as assuming *whatever is necessary* to convert the kinds of stimuli we're exposed to into experiences with the determinate contents that we know our experiences have by virtue of knowing their phenomenal characters.

On option (a), sub-personal states of our perceptual systems are genuinely contentful (as opposed to having content metaphorically ascribed to them), but their contents are *indeterminate*, owing to the fact that there are numerous different ways of characterizing the regularities that our perceptual systems have adapted to pick up on, which the mechanisms of selection are indifferent between (Dennett 1987, chap.8; Shea 2018, 148-62). Consequently, there will be no determinate fact of the matter as to what *precisely* our perceptual systems assume if viewed solely in this way. We'll have no way of deciding, e.g., whether they have adapted to treat certain spectral discontinuities as indicating the presence of edges, or edge stages, or the disjunction of all such gawagai-ish edge-variants. Our inability to provide a determinate characterization of the content of these assumptions when viewed in this manner is the reason why they are unable to account for the determinate contents of our experiences on their own. Hence, again, the Structuralist's insistence that the content of our experiences is instead determined by their phenomenal character.

On option (b), the assumptions built into our perceptual systems *would be* sufficient to determine the contents of our experiences, but this is only because we've now tailored our characterization of those assumptions to ensure this, by treating them as assuming whatever is necessary to produce experiences with the determinate contents that we already know our experiences have by virtue of knowing their phenomenal characters. In short, we start from a position of knowing what our experiences represent simply by virtue of knowing what it's like to have them. Any ascriptions of *determinate* content to the sub-personal states involved in producing our experiences are then guided and made possible by this knowledge, by adopting Dennett's

(1987) ‘design stance’ towards our own neural apparatus and considering the roles that different parts of it play in producing experiences whose content is determined by (and known to us by way of) their phenomenal character. This enables us to ascribe determinate content to the assumptions that our perceptual systems implement, but only at the cost of making these ascriptions derivative and non-literal, since we are now using the phenomenally determined content of the experiences they produce to decide among various potential characterizations of their content that the mechanisms of selection are indifferent between.

Distinguishing between these two ways of viewing the assumptions built into our perceptual systems helps resolve the objection raised above. The worry, again, is that these assumptions are supposed to help explain why our experiences are grouped in the way that they are and thus have the contents they do, yet the nature of these assumptions is in turn said to depend on the phenomenally determined content of experiences produced on their basis. In this statement, however, the assumptions are being conceived in two different ways. The first clause is true only if they are viewed *à la* option (a), whereas the second is true only on option (b). On the latter option, the assumptions built into our perceptual systems indeed cannot explain experiential content, since our characterization of those assumptions is (on this construal) non-literal and itself based on our prior knowledge of the contents of our experiences. On option (a), however, the content we ascribe to our perceptual systems in treating them as relying on certain assumptions is indeterminate but literal, and not derived from the phenomenally determined content of the experiences produced on their basis. On this construal, the assumptions built into our perceptual systems thus *can* help to explain experiential content, by explaining why certain relations among experiences enable them to be grouped and externalized while others do not, and why experiences whose components stand in certain basic structural relations are thereby grouped in such a way as

to represent the general *kinds* of features that they do (e.g. edges or some gavagai-ish variant thereof, instead of some completely different feature, e.g. weight, or pitch). Viewed in this way, these assumptions cannot fully account for the determinate contents of our experiences, since on this conception, the assumptions are themselves indeterminate. This, however, is precisely why PITists treat an experience's phenomenal character as having the final say in what its content ultimately is. Contrary to the foregoing objection, Structuralists can thus reasonably hold that the assumptions built into our perceptual systems help explain why our experiences are organized in the way that they are and thus have the contents that they do even if an experience's content is ultimately determined by its phenomenal character alone.

6. Conclusion

This paper has defended Structuralist PIT against two objections: (i) that it cannot explain why some experiences have content while others don't, and (ii) that it cannot explain why contentful experiences have the specific contents that they have. In response to (i), I've argued that to possess content, an experience must present itself as an experience of something distinct from itself, and only experiences that are both sufficiently complex and organized are capable of doing this, as only such experiences present themselves as being not random or self-generated, but instead as having been produced by things distinct from themselves, which they thereby inform us of. Elaborating on this idea, I've further suggested that the relations that enable our experiences to be externalized in this way involve (a) high degrees of qualitative similarity among spatially and temporally contiguous experiences, (b) smooth contours among similarity-based groupings, and (c) cross-modal agreement in similarity-based groupings. In response to (ii), I've argued that while all forms of PIT, Structuralism included, must treat the connection between an experience's

phenomenal character and its content as to some extent brute, Structuralists can nevertheless provide a deeper explanation of perceptual content than other forms of PIT by showing how the content-determining groupings among our experiences depend on the assumptions that our perceptual systems have adapted to use in processing sensory input. On the resulting picture, the content of our experiences depends ultimately on their phenomenal character, but our biological history also helps to explain why our experiences have the phenomenal characters (and thus the contents) that they do by explaining why their components are organized in the way that they are.¹³

¹³ Many thanks to two anonymous reviewers for their helpful comments.

References

- Adams, Frederick, and Kenneth Aizawa. 1994. 'X' Means X: Fodor/Warfield Semantics. *Minds and Machines* 4: 215–31.
- Baker, Lynn Rudder. 1997. Why Constitution is Not Identity. *Journal of Philosophy* 94: 599–621.
- Biro, John. 2018. Constitution and Identity. *Erkenntnis* 83: 1127–38.
- Blakemore, Colin, and Grahame Cooper. 1970. Development of the Brain depends on the Visual Environment. *Nature* 228: 477–8.
- Bourget, David. 2010. Consciousness Is Underived Intentionality. *Noûs* 44: 32–58.
- Bregman, Albert. 1993. Auditory Scene Analysis: Hearing in Complex Environments. In *Thinking in Sound: The Cognitive Psychology of Human Audition*, eds. McAdams and Bigand. Oxford University Press.
- Casati, Roberto. 2015. Object Perception. In *The Oxford Handbook of Philosophy of Perception*, ed. Matthen. Oxford University Press.
- Chalmers, David. 2006. Perception and the Fall from Eden. In *Perceptual Experience*, eds. Gendler and Hawthorne. Oxford University Press.
- Delk, John, and Samuel Fillenbaum. Differences in Perceived Color as a Function of Characteristic Color. *American Journal of Psychology* 78: 290–3.
- Dennett, Daniel. 1987. *The Intentional Stance*. MIT Press.
- Dretske, Fred. 1988. *Explaining Behavior*. MIT Press.
- Farkas, Katalin. 2013. Constructing a World for the Senses. In *Phenomenal Intentionality*, ed. Kriegel. Oxford University Press.
- Fine, Kit. 2003. The Non-Identity of a Material Thing and Its Matter. *Mind* 112: 195–234.
- Fodor, Jerry. 1987. *Psychosemantics*. MIT Press.
- Gates, Gary. 1996. The Price of Information. *Synthese* 107: 325–47.
- Gibson, James. 1979. *The Ecological Approach to Visual Perception*. Houghton Mifflin.
- Hall, Richard. 1990. Does Representational Content Arise from Biological Function? *Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1: 193–99.
- Hansen, Thorsten, Maria Olkkonen, Sebastian Walter, and Karl Gegenfurtner. 2006. Memory Modulates Color Appearance. *Nature Neuroscience* 9: 1367–8.
- Horgan, Terence, and John Tienson. 2002. The Intentionality of Phenomenology and the Phenomenology of Intentionality. In *Philosophy of Mind: Classical and Contemporary Readings*, ed. Chalmers. Oxford University Press.
- Horgan, Terence, and George Graham. 2012. Phenomenal Intentionality and Content Determinacy. In *Prospects for Meaning*, ed. Schantz. De Gruyter.
- Johnston, Mark. 1992. Constitution is Not Identity. *Mind* 101: 89–105.
- Kriegel, Uriah. 2013. The Phenomenal Intentionality Research Program. In *Phenomenal Intentionality*, ed. Kriegel. Oxford University Press.
- Lewis, David. 1971. Counterparts of Persons and Their Bodies. *Journal of Philosophy* 68: 203–11.
- Loar, Brian. 2003. Phenomenal Intentionality as the Basis of Mental Content. In *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, eds. Hahn and Ramberg. MIT Press.
- Masrour, Farid. 2013. Phenomenal Objectivity and Phenomenal Intentionality: In Defense of a Kantian Account. In *Phenomenal Intentionality*, ed. Kriegel. Oxford University Press.
- Mendelovici, Angela. 2018. *The Phenomenal Basis of Intentionality*. Oxford University Press.
- Millikan, Ruth. 1989. Biosemantics. *Journal of Philosophy* 86: 281–97.
- Neander, Karen. 2017. *A Mark of the Mental: In Defense of Informational Teleosemantics*. MIT Press.

- Noonan, Harold. 1993. Constitution is Identity. *Mind* 102: 133–46.
- O’Callaghan, Casey. 2012. Perception and Multimodality. In *The Oxford Handbook of Philosophy of Cognitive Science*, eds. Margolis, Samuels, and Stich. Oxford University Press.
- Palmer, Stephen. 1999. *Vision Science: Photons to Phenomenology*. MIT Press.
- . 2003. Visual Perception of Objects. In *Handbook of Psychology, Volume 4: Experimental Psychology*, eds. Healy and Proctor. Wiley.
- Quine, W. V. 1960. *Word and Object*. MIT Press.
- Shea, Nicholas. 2018. *Representation in Cognitive Science*. Oxford University Press.
- Searle, John. 1992. *The Rediscovery of the Mind*. MIT Press.
- Spelke, Elizabeth. 1990. Principles of Object Perception. *Cognitive Science* 14: 29–56.
- Strawson, Galen. 2010. *Mental Reality*. 2nd ed. MIT Press.
- Woodward, Philip. 2019. Phenomenal Intentionality: Reductionism Vs. Primitivism. *Canadian Journal of Philosophy* 49: 606–27.