

Rainbow's end: The structure, character, and content of conscious experience

Brandon Ashby 

Philosophy Department, University of Arizona, Tucson, Arizona

Correspondence

Brandon Ashby, Philosophy Department, University of Arizona, P.O. Box 210027, Tucson, AZ 85721.
Email: bashby.philosophy@gmail.com

Separatism, representationalism, and phenomenal intentionalism are the primary views on the relationship between the phenomenality and intentionality of experience. I defend a novel position that is incompatible with separatism, can enrich representationalism and phenomenal intentionalism, but can also be accepted independently of those views. I call it *phenomenal schematics*: The phenomenal characters of our experiences have structures that place a priori, formal, and sometimes semantic constraints on our experience's possible intentional contents. Phenomenal structures are like the grammar of a language (or the compositional rules governing maps, models, and diagrams). Unlike words, however, phenomenal characters possess their "grammatical properties" essentially.

KEYWORDS

experiential content, phenomenal intentionality, phenomenal schematics, phenomenal structure, representationalism, separatism

1 | INTRODUCTION

Consider the difference between pressure experiences and temperature experiences. What it is like to feel a pressure has two components: One corresponding to the degree of force and the other to the direction of force. If you press your left index finger straight down into the palm of your right hand, and then at a 45° angle, you can match the two experiences in terms of the felt degree of force while each differs in the felt direction of force. To formally model the phenomenal characters of our pressure experiences, we would need to use vectors (or vector fields).¹ They have a vectorial structure. Temperature experiences, however, have no second

¹Phenomenal characters are the qualitative aspects of our experiences, what it is like to have them.

component. We can be more or less hot or cold, cool or warm. There is nothing analogous to a direction of force component. To formally model our temperature experiences, we would need to use directionless magnitudes (or scalar fields). They have a scalar structure.

I shall argue that structural features of the phenomenal characters of our experiences, like the above examples, place formal and in some cases semantic constraints on the possible intentional contents that our experiences can have, and these constraints hold with a priori necessity. I call this view *phenomenal schematics*.² For instance, consider the phenomenal character of our experiences of warmth, *p-warmth*. I shall argue that the structure of *p-warmth* rules out the possibility that it could have served to represent a pressure; *p-warmth* lacks a second component that could serve as a means of representing a direction of force. Not all possible contents besides *warmth* will be ruled out by the phenomenal structure of *p-warmth*, however; phenomenal schematics is compatible with (though it does not entail) the claim that *p-warmth* could have represented coolness instead of warmth.

There are currently three main views on the relationship between phenomenality and intentionality: Separatism, representationalism, and phenomenal intentionality. Phenomenal schematics is incompatible with separatism, but it is compatible with and can enrich representationalism and phenomenal intentionalism. Nevertheless, phenomenal schematics can be accepted on its own terms and does not require accepting either representationalism or phenomenal intentionalism. Phenomenal schematics marks an overlooked (non-terminal) node in dialectical space (see Figure 1).³

Separatism claims that phenomenality and intentionality are contingently related. As David Papineau (2016) says while defending the view, separatism claims that phenomenal characters, “in themselves are like typographical words, items that have no constitutive tie to what they contingently represent” (p. 343). Separatism was the default position in analytic philosophy of mind until the 1980s (Lycan, 2008). While few philosophers explicitly defend the view today—Jaegwon Kim (2005) is widely interpreted as a separatist, and David Papineau (2014, 2016) is the view's most recent advocate—separatism remains popular, nonetheless. Papineau (2014) even claims that separatism is, “little more than common sense” (p. 2).

Representationalism claims that phenomenality is identical to, reducible to, or grounded in a kind of intentional content (Byrne & Tye, 2006; Chalmers, 2006; Dretske, 1996; Lycan, 2001). In short, there is a kind of intentionality that suffices for phenomenality with metaphysical necessity.

Phenomenal intentionalism claims that there is a kind of intentionality that is identical to, reducible to, or grounded in phenomenality (Horgan & Tienson, 2002; Kriegel, 2011; Loar, 2003; Searle, 1992; Strawson, 2008). In short, phenomenality suffices for a kind of intentionality with metaphysical necessity.

We can understand the relationships between phenomenal schematics, separatism, representationalism, and phenomenal intentionalism in terms of their answers to three questions:

- (1) Are there any a priori or metaphysically necessary constraints between the phenomenal characters and intentional contents of our mental states?

²I remain neutral on whether these constraints are metaphysically necessary, or whether their a priori status is due to contingent features of our cognitive architecture. Even if the latter obtains, the constraints identified below will plausibly hold for all (non-pathological) human experience (see fn. 13).

³I shall not discuss disjunctivism or naïve realism, which deny that our experiences have contents. My arguments are addressed to realists about experiential contents.

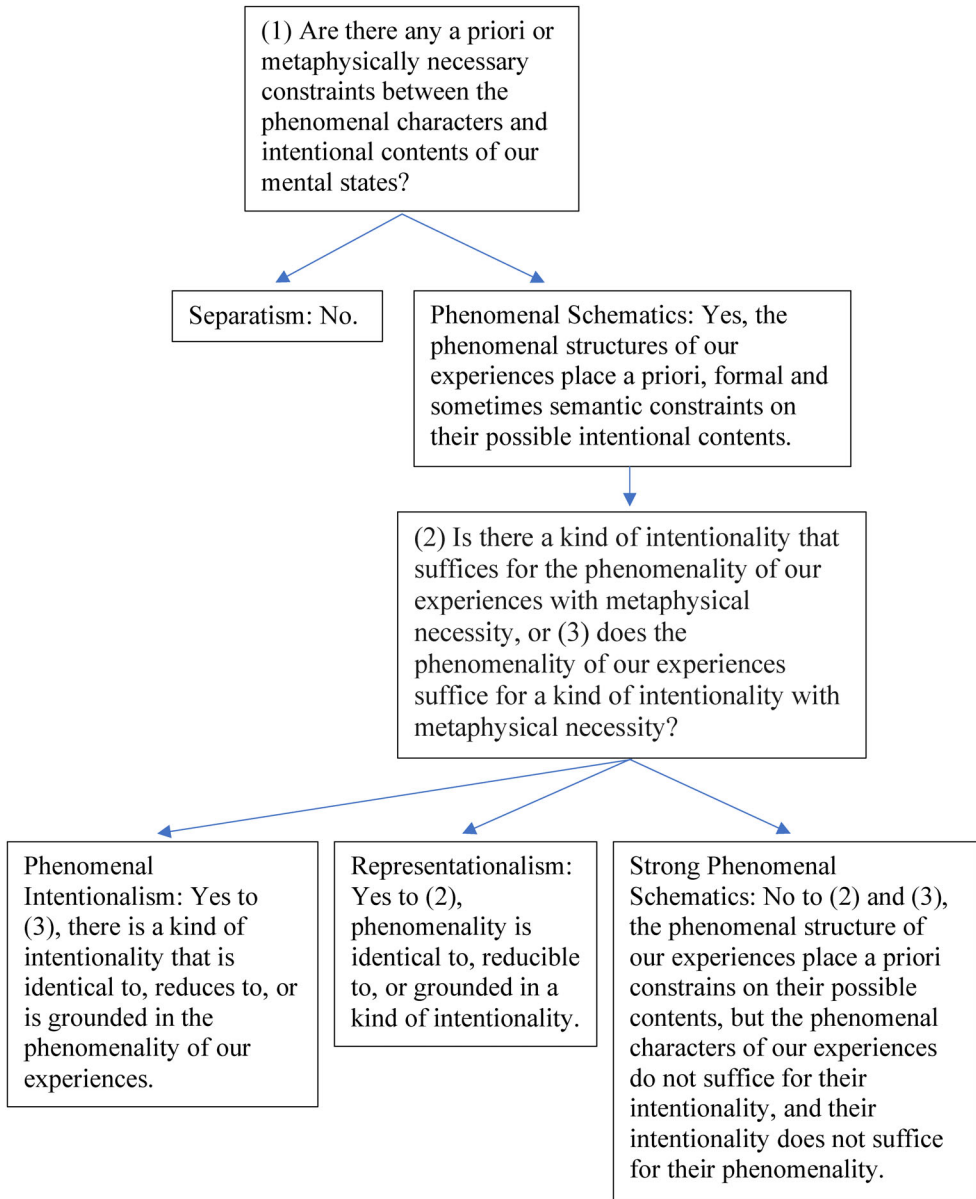


FIGURE 1 The shape of the debate over the relationship between the intentionality and phenomenality of our experiences

Separatism answers (1) in the negative; phenomenal schematics, representationalism, and phenomenal intentionalism all answer in the affirmative.

(2) Is there a kind of intentionality that suffices for the phenomenality of our experiences with metaphysical necessity?

(3) Does the phenomenality of our experiences suffice for a kind of intentionality with metaphysical necessity?

Representationalists and phenomenal intentionalists respectively answer (2) and (3) in the affirmative.⁴ Phenomenal schematics, however, remains *neutral* on how to answer (2) and (3). It makes no claims about whether or not phenomenality ever suffices for intentionality with metaphysical necessity, or vice versa. Phenomenal schematics is distinct from—but compatible with—representationalism and phenomenal intentionalism.

My aims are, first, to demonstrate that separatism is false while remaining neutral on the status of representationalism and phenomenal intentionalism, and, second, to offer a detailed characterization of the phenomenal structure of our experiences and the specific kinds of constraints that it places upon their possible contents. The result is a novel perspective on the nature of phenomenal consciousness that is available to all non-separatist philosophers of mind.

I begin by presenting my argument for phenomenal schematics (Section 2) and addressing a possible objection (Section 3). I then compare my own method for exploring phenomenal structure with prior efforts (Section 4). Next, I show that Kim's and Papineau's arguments for separatism are not strong enough to establish the view (Section 5). And I conclude with phenomenal schematics' central lesson: The phenomenal structures of our experiences are akin to the grammatical properties of written words (or the rules of composition governing the representational elements in diagrams, maps, and models); unlike words, however, phenomenal characters possess their "grammatical properties" essentially.

2 | THE ARGUMENT FOR PHENOMENAL SCHEMATICS

I appeal to a range of variations on the inverted color spectrum and zombie thought experiments in order to motivate phenomenal schematics. The former sort of case involves *phenomenal displacement subjects*, perfect physical/functional duplicates of ourselves whose experiences of a given kind map onto our twins' physiology differently than they map onto our own. Temperature spectrum inverts are one example: What it is like for them to feel moderately warm things is what it is like for us to feel moderately cool things, and vice versa, and so on for other temperatures. The latter sort of case involves *partial zombies*: Physical/functional duplicates of us who lack some but not all aspects of what it is like to be us. A *visual zombie*, for instance, lacks any and all visual experiences, but is otherwise just like us, including her non-visual experiences and her physical/functional makeup.⁵

As we shall see, some displacement cases turn out to be *inconceivable* in the sense that they are a priori incoherent: They ask us to consider experiences with incoherent phenomenal structures. Some partial zombie cases turn out to be *unimaginable* in the sense that we cannot imagine what the partial zombie's experiences are like (even though those experiences are plausibly a priori coherent).⁶ These thought experiments ask us to consider experiences with alien

⁴Philosophers who claim that phenomenality is identical to a kind of intentionality (e.g., Chalmers, 2006) answer "yes" to (2) and (3). They qualify as both representationalists and phenomenal intentionalists.

⁵Subjects with sensory deficits, such as blind subjects, are not partial zombies. Partial zombies are physical/functional duplicates of actual subjects who have the kinds of experiences that the partial zombie lacks. So, a visual zombie will be a physical/functional duplicate of a normally sighted person, and blind subjects are not physical/functional duplicates of sighted subjects.

⁶The unimaginability of partial zombies differs from extant examples of unimaginable experiences, such as, bat experiences (Nagel, 1974). In the case of unimaginable partial zombies, we are already familiar with *every aspect* of what it is like to be the partial zombie; partial zombies *only* differ from us in *lacking* some aspects of what it is like to be us. There are no unfamiliar phenomenal characters involved, unlike the bat case.

phenomenal structures. And finally, the conceivable displacement and imaginable partial zombie cases ask us to consider experiences with phenomenal structures that are identical to our own. What these cases will collectively demonstrate is that the phenomenal structure of our experiences is a compositional and systematic structure akin to the grammar of a language (or the rules for combining different representational elements in diagrams, maps, and models). Phenomenal structure governs how different kinds of phenomenal characters can and cannot be combined and thereby constrains our experiences' possible contents.⁷

2.1 | Displacement cases

Let us begin with a handful of conceivable displacement cases. The inverted color spectrum is one example, and the inverted temperature spectrum is another. Olfactory and gustatory displacement cases also seem to be conceivable. It seems relatively easy to imagine subjects for whom what it is like to taste sweet things is what it is like for us to taste sour things, and vice versa, or for whom ammonia smells like roses, and vice versa.⁸

2.1.1 | Visual-shape displacement

Let us now consider some a priori incoherent cases. Try to imagine perfect physical/functional duplicates of ourselves for whom what it is like to see spheres is what is like for us to see cubes, and vice versa. This *visual-shape displacement case* can be deeply puzzling. For instance, what is it like for our visual-shape displaced twins to see other kinds of shapes, such as pyramids, ellipsoids (rugby balls), or tori (donuts)? What about the shadows and cross-sections of cubes and spheres? For normal subjects, what it is like to see a sphere is the same regardless of the sphere's orientation; not so for cubes. So, do rotating spheres look like rotating cubes or stationary cubes to our twins? And what determines whether a sphere looks like a cube seen from one orientation as opposed to another?

These questions are difficult to answer, and there is a plausible reason why. The color and temperature spectrum inversion cases provide us with a general procedure for translating between our experiences and our twins' experiences. We just "find and replace" red experiences with green experiences, and vice versa, and so on. In doing so, we face no conflicts, where we must replace a red experience with a green experience *and also* a purple experience. And we face no ambiguities, where we could replace a red experience with *either* a green experience *or* a blue experience. In contrast, the visual-shape displacement case gives us no general procedure for translating between our own and our twins' experiences of shape; indeed, it could not. This is because (a) cubes have geometrically privileged parts and spheres do not, (b) the phenomenal characters of our experiences of the parts of cubes and the parts of spheres reflect their

⁷Appeals to inverts, zombies, and variations on these cases are common in the debates over the phenomenal-intentional relation (e.g., Chalmers, 2006; Horgan, 2014; Kim, 2005; Papineau, 2014, 2016). In claiming that some but not all displacement and partial zombie cases are conceivable/imaginable, I commit myself to the existence of an explanatory gap, which a small minority of physicalists will object to. I am no more committed on this front, however, than is common in the present debate.

⁸These cases may disrupt the patterns of similarity between our experiences of different odors and tastes, such as the similarities between the scents of different kinds of flowers.

respective heterogeneity and homogeneity, and (c) our experiences have a systematic and compositional character to them. Allow me to elaborate.

The phenomenal characters of our visual experiences are *systematic* and *compositional*: If we hold orientation, color, and so forth fixed, then what it is like to see a three-dimensional, right-angled vertex as part of a cube is largely the same as what it is like to see a three-dimensional, right-angled vertex as the peak of a three-sided pyramid. Of course, we may need to accommodate gestalt, repetition, fatigue, and attentional effects as well as the differences between central and peripheral vision. But even when we take such influences into account, there remains a striking degree of similarity between our experiences of one three-dimensional, right-angled vertex and another regardless of what sort of shape that vertex is seen as a part of. Moreover, what it is like for us to see a red cube, for instance, is a function of what it is like for us in general to see red things and what it is like for us in general to see cubes. The phenomenal characters of our total experiences at a time are (typically) a function of the less-than-total phenomenal characters that make them up.⁹

Visual-shape displacement necessarily violates the systematic and compositional character of our experiences. Whereas cubes have geometrically privileged parts—straight edges, square faces, and right-angled vertices—spheres lack geometrically privileged parts—they only have regions of positive, uniform curvature. Importantly, our experiences of the parts of cubes reflect their heterogeneity: What it is like to see a vertex of a cube differs from what it is like to see an edge of a cube. And our experiences of the various surface regions of spheres reflect their homogeneity: If we view different surface regions of a sphere from the same orientation, and we hold their color and texture fixed, one part of a sphere looks like any other.

So, we would need a *many-to-one mapping* to “translate” from our visual experiences of cubes and their parts to our twins’ experiences of cubes and their parts. Our twins’ experiences of what are in fact straight edges, square faces, and three-dimensional, right-angled vertices would all need to have the phenomenal character that our experiences have when we see regions of positive, uniform curvature. And we would need a *one-to-many mapping* to translate from our visual experiences of spheres and their parts to our twins’ experiences of spheres and their parts. Some but not all of our twins’ experiences of what are in fact regions of positive, uniform curvature would need to have the sort of phenomenal character that our experiences have when we look at three-dimensional, right-angled vertices, and similarly for straight edges and square faces.

Moreover, our experiences of the parts of spheres and cubes can be parts of our experiences of other shapes.¹⁰ Hemispheres have regions of three-dimensional, positive, uniform curvature. And any number of shapes besides cubes have straight edges, square faces, and three-dimensional, right-angled vertices. The many-to-one and one-to-many mappings between our experiences and our twins’ experiences of the parts of cubes and spheres prevent us from being able to identify a unique phenomenal character that would be associated with our twins’ experiences of straight edges—and similarly for square faces, right-angled vertices, and regions of positive, uniform curvature. Any efforts to enforce the systematicity and compositionality of our experiences would terminate in conflicting (many-to-one) and ambiguous (one-to-many) displacements.

Visual-shape displacement violates the systematicity and compositionality of our experiences and, consequently, undermines the very features of our visual experiences of shape that enable them to systematically represent the world around us. Our experiences of right-angled vertices, for instance, can serve to systematically represent right-angled vertices because our

⁹Gestalt experiences may be an exception.

¹⁰Bayne and Chalmers (2003) also attribute a (quasi)mereological structure to consciousness.

shape experiences that include right-angled vertices as geometrically privileged parts are themselves partly composed out of our experiences of right-angled vertices.¹¹ For our twins, there could be no one phenomenal character that could serve in general to represent right-angled vertices and that could always play the same role in composing their experiences of shapes that contain right-angled vertices. Our twins' shape experiences would lack the systematic and compositional structure that allows those experiences to function as a means of representing shapes, their parts, and arrangements of those parts.

There are additional problems with visual-shape displacement. The phenomenal characters of our experiences of cubes and spheres are systematically related to the phenomenal characters of our experiences of other shapes even when those other shapes contain none of the geometrically privileged parts of cubes or spheres. Ellipsoids contain no regions of uniform, three-dimensional curvature, yet the closer an ellipsoid approximates perfect sphericity, the more similar our experiences of that ellipsoid will be to our experiences of spheres. More generally, similar shapes tend to produce experiences with similar phenomenal characters. Visual-shape displacement would necessarily violate this feature of our experiences due to the requirement for many-to-one and one-to-many mappings between our twins' shape experiences and our own.

At least some projective-geometric relationships seem to be latent in visual experience. Consider what it is like to look at a cube that is casting a shadow onto a warped tabletop. The shapes of the cube, its shadow, and the tabletop may not have any geometrically distinguished parts in common. Nevertheless, the shapes of the cube, its shadow, and the tabletop are all systematically related to one another, and these relationships appear to be reflected in the phenomenal characters of our visual experiences. Our experience of the shape of the shadow "makes sense" given our experiences of the shapes of the cube and the tabletop. Subjects unfamiliar with projective geometry would struggle to articulate what this system of relationships is, but it seems plausible that this network of relationships will often be isomorphic to those described by a projective, approximately Euclidean geometry.¹² Visual-shape displacement disrupts this feature of our visual experiences.

What these observations suggest is that our visual experiences possess not only a compositional and systematic structure, but a *geometric* structure as well. The geometric relationships that are definable between the shapes of objects and their parts are themselves reflected in the relationships between our experiences of those shapes and their parts.¹³ Just as squares are composed of straight edges, right-angles, and their arrangement, our experiences of squares are composed of our experiences of straight edges, right-angles, and their arrangement. Visual-shape displacement undermines the geometric structure of our visual experiences by requiring many-to-one and one-to-many mappings between our twins' and our own experiences of the parts of cubes and spheres.

2.1.2 | **Yellowness-itchiness displacement**

Try to imagine a subject for whom what it is like to feel itchy is what it is like for us to see yellow, and vice versa. The widespread awareness of synesthesia can lead to misinterpretations of

¹¹Michael Tye (2003) claims that only total experiences exist; we can reformulate the present discussion in terms of ways that total experiences are similar to or different from one another.

¹²Visual experience's representation of space may violate the parallel postulate (Masrour, 2014). Indeed, there may be no single geometry for visual space (Wagner, 2006). However, the non-Euclidean features of our visual experiences are quite subtle and require careful experimentation to detect.

¹³Visual form agnosia subjects' experiences (Farah, 2004) plausibly lack this structure.

the case that I want to avoid.¹⁴ I am asking us to imagine a subject whose experiences of itchiness play the exact same role in composing her visual experiences that our experiences of yellow do in composing our visual experiences. Her experiences of itchiness combine with her visual experiences of shape, size, distance, texture, and colors besides yellow in order to compose her visual experience of, say, a blue to yellow color gradient.

Moreover, I am asking us to imagine a subject whose experiences of yellow play the exact same role in composing her experiences of her body surface that experiences of itchiness play in shaping our experiences of our body surfaces. Experiences of yellow combine with experiences of warmth, felt-texture, and pressure to compose the sort of experience that someone who is allergic to cats might have when a cat rubs against her.

The case is odd, to say the least. And the reason seems to be that experiences of yellowness and experiences of itchiness play very different sorts of roles from one another in composing what it is like to be us at a time. One key difference stems from the fact that visual and bodily experiences appear to utilize *different kinds of representation of space*. Vision appears to use a three-dimensional, projective, and approximately Euclidean representation of space (Section 2.1.1). The senses of touch, proprioception, and kinesthesia seem to employ a topological representation of the body-space. Itchiness is represented as a kind of state that regions of the body surface can take on, one that compels us to scratch the affected area until the state dissipates. The joints are represented as points of articulation to which forces are applied endogenously through muscle activity and exogenously by external objects. Temperatures occupy areas of the body surface or regions of the body volume. And pressures are experienced as deforming the resting shape of the body and as requiring certain kinds of actions to resist or overcome.

Audition appears to use a radial representation of space that concerns the angle of the sound source with respect to the subject's head and whether the sound source is approaching or receding, carrying little to no information about the shapes, sizes, and distances of heard objects.

Olfaction and gustation appear to carry no spatial information at all. Smells and tastes have locations, but our olfactory and gustatory experiences do not themselves represent a space in which those odors and tastes are located; other senses carry that information.

My suggestion is that displacing experiences across sensory modalities will generally prove inconceivable because our experiences of different sorts of properties are "*formatted*" in terms of the different kinds of spatial representations that our sensory modalities employ. The ways in which our less-than-total experiences within a given modality combine to compose our total experiences within that modality are defined in terms of how they can be located in whatever representation of space that modality employs. After all, the idea that what it is like for us to *feel* squares could have been what it is like for us to *see* squares, and vice versa, hardly seems any more comprehensible than the yellowness-itchiness displacement case. Bodily experiences of itchiness and felt squares are formatted in terms of a topological representation of the body-space; visual experiences of yellowness and seen-squares are formatted in terms of a projective and approximately Euclidean representation of visible-space.

¹⁴I am not asking us to imagine a subject who itches whenever she sees something yellow or who has a sensory image of yellow in her visual field whenever she feels itchy. And I am not asking us to imagine a subject who has visual illusions that parts of her bodily surface turn yellow when they itch, or who sees yellow objects as suffering from itchiness.

2.1.3 | Pressure–temperature displacement

Let us return to our introductory example and consider physical/functional duplicates of ourselves for whom what it is like to feel temperatures is what it is like for us to feel pressures, and vice versa. If you press your left index finger straight into your right palm, and then at a 45° angle, you can match the two experiences in terms of the felt degree of force even though they differ in the felt direction of force. So, there are two components to our experiences of pressure: One corresponding to the degree of force and the other to the direction of force. To formally model our pressure experiences, we shall need to use *vectors* (or vector fields). Temperatures, however, only come in degrees. They have no “second component.” To formally model our temperature experiences, we shall need to use *directionless magnitudes* (or scalar fields). Positive numbers could represent the warm/hot side of the spectrum, negative numbers could represent the cool/cold side, and zero could represent body-neutral temperatures.¹⁵

The pressure–temperature displacement case provides us with our first glimpse of how we could formally characterize phenomenal structure. Different sorts of experiences will require different sorts of formal representations—vector fields versus scalar fields—and experiences that require different sorts of formal representations are not displaceable even when they both belong to a single sensory modality or cluster of sensory modalities that involve a common spatial formatting—such as the haptic, kinesthetic, and proprioceptive modalities.

To emphasize the point, consider itchiness again. We can be more or less itchy; our experiences of itchiness are best formally represented by directionless magnitudes (or scalar fields). But, unlike our experiences of warmth, there is no clear opposite of our experiences of itchiness. A good formal characterization of our experiences of itchiness would use zero to represent no itchiness and ever larger numbers to represent increasing degrees of itchiness. Formal representations of temperature experiences, in contrast, will require an axis of symmetry around which distinct but symmetrical experiences of hot and cold, cool and warm are located.

2.2 | Partial zombies

Let us consider a handful of imaginable partial zombies before moving on to some unimaginable (but a priori coherent) cases. Examples of imaginable partial zombies are plentiful. For instance, it seems easy to imagine partial zombies who lack the experiences associated with any particular sensory modality; visual zombies, auditory zombies, and so forth. It also seems fairly easy to imagine zombies who lack any and all experiences associated with a given region of egocentric space. A partial zombie might lack all experiences associated with the right half of her visual field, the left half of her body, or anything that would be heard as coming from behind her.

What the imaginability of these cases suggests is that we think of sensory modalities and regions of egocentric space as forming “partitions” such that we can imagine eliminating all the experiences associated with one partition without rendering the experiences associated with other partitions unrecognizable. For example, we can imagine zombifying our visual experiences without being forced to imagine dramatic changes in our auditory or haptic experiences that would render them unrecognizable as auditory or haptic experiences of the sort had by

¹⁵A more complex model will ultimately be required: Experiences of increasingly hot and cold temperatures become increasingly indistinguishable.

healthy, adult subjects. And we can imagine zombifying our experiences of the right halves of our bodies without having to imagine any dramatic differences in our experiences of the left halves of our bodies that would render them unrecognizable as the sorts of experiences had by normal subjects.¹⁶

We cannot always zombify a kind of experience without radically altering the remaining experiences. Consider the *visual-shape zombie*. She has no visual experiences of shape whatsoever, but she does have visual experiences of size, motion, orientation, texture, color, and luminance. What it is like for her to see a circle does not differ from what it is like to see a square so long as the two objects are matched for area, color, and so forth. The visual-shape zombie's visual experiences are as shapeless as the taste of coffee.

The case is hard to imagine (though not incoherent). And the reason seems to be that it asks us to consider experiences where the structure of the zombie's residual visual experiences differs from the structure of those sorts of experiences when had by normal subjects. The visual-shape zombie's visual experiences lack the geometrical structure that ours possess (Section 2.1.1). And this is a structure that our visual experiences of properties besides shape require in order to play their normal role in composing our total visual experiences. We do not see the size of an object independently of its shape; we see that size as the size of that object's shape. And when we see the color and the texture of an object, its color and texture "fill out" its shape. We cannot eliminate our visual experiences of shape without altering how our color, texture, and size experiences combine to compose our total visual experiences in the way that we could eliminate the predicate *square* from $\exists x(\text{red}(x) \ \& \ \text{large}(x) \ \& \ \text{square}(x))$ without altering the role that the other predicates play in composing that sentence. The systematic roles that our experiences of shape, color, texture, and size play in composing our total visual experiences are interdependent.

Finally, consider the *color-luminance zombie*: A subject for whom there is nothing that it is like to see different colors, different degrees of illumination (sunlight vs. moonlight), and different shades of black and white and gray. The color-luminance zombie only has visual experiences of shape, size, texture, orientation, distance, and motion. For this subject, there is no difference in what it is like to see a black or a white cube, a brightly lit or a dimly lit cube, or a red or a green cube so long as they are otherwise identical. Her visual experiences are *purely spatial*. Even in the case of transparent objects, the color-luminance zombie does not experience any alterations (such as magnification or warping) of the color-luminance boundaries of the objects that are seen through those objects. Her visual experiences are as colorless and luminance-free as the sound of breaking glass.

What the visual-shape and color-luminance zombies conjointly suggest is that the systematic role that our experiences of color play in composing our total visual experiences rely upon our experiences of shape playing their role in composing our total visual experiences, and vice versa. More generally, the roles that our experiences of primary and secondary qualities play in composing our total experiences within a given modality appear to be mutually interdependent. Secondary qualities are formatted in terms of the kind of spatial representation that they are embedded in (Section 2.1.2). And primary properties, at least for us, appear to require secondary qualities to "fill them out."¹⁷

¹⁶I shall remain neutral on the possibility of multi-sensory experiences. And, I do not deny that there are gestalt experiences that range over regions of the visual field or the body surface. My claim is that we can imagine eliminating the visual component and any multi-sensory components from, say, an audio-visual experience of a popping balloon and retain an auditory experience that is recognizable as a relatively normal auditory experience. Moreover, my claim is that we can imagine zombifying the right half of our visual fields such that, when we look at duck-rabbit drawings, our residual experiences in the left halves of our visual fields are still recognizable as fairly ordinary visual experiences.

¹⁷There is much more to be said on this topic than can be said here.

2.3 | Bringing it all together

Our intuitions concerning displacement and partial zombie cases are quite systematic. In the conceivable displacement cases, the displaced experiences play equivalent roles in composing our total experiences at a time. In the readily imaginable partial zombie cases, we imagine subtracting experiences in such a way that the structure of the residual experiences remains unchanged.

Unimaginable partial zombie cases ask us to consider experiences with alien phenomenal structures: The role that the residual experiences play in composing our total experiences depend upon the role that the zombified experiences play in composing our total experiences. Consequently, the residual experiences cannot play their normal role in composing our total experiences at a time.

A priori incoherent displacement cases ask us to imagine experiences with incoherent phenomenal structures. They violate the compositional and systematic structure that our experiences have and in virtue of which any given less-than-total experience, such as our visual experience of a straight line, can serve to systematically represent a particular subject matter, such as straight lines.

While we shall return to these issues in Section 5, what we see is that phenomenal structures are akin to the grammatical properties of words or the rules of composition governing the representational elements employed in a diagrams, maps, or models (such as, topographic lines). Just as one-place predicates can only represent properties and two-place predicates can only represent binary relations, temperature experiences can only represent scalar magnitudes that act on the body surface or that are distributed throughout the body volume; temperature experiences have a scalar structure defined over a topological representation of the body-space. Moreover, just as we could not use images of squares to represent circles and images of circles to represent squares when drawing geometric figures in Cartesian coordinate planes, we could not displace the phenomenal characters of our experiences of cubes with the phenomenal characters of our experiences of spheres. Doing so would violate the systematic, compositional, and geometric structure of our visual experiences. Even if the specific phenomenal characters of our experiences do not uniquely determine the specific contents of those experiences—I neither affirm nor deny that claim here—our experiences possess different kinds of phenomenal structure that constrain our experience's possible contents in different ways.¹⁸

3 | REPRESENTATION IN PERCEPTUAL EXPERIENCE

There are different theories of representation. One theory in particular makes it unlikely that the characteristics of a representation—be it an experience, a sentence, or what have you—can

¹⁸In tactile-visual sensory substitution (TVSS), a 20 x 20 grid of vibrating/electrically conductive nodes is placed somewhere on a subject's body. The nodes are connected to a camera and stimulate the skin with an intensity proportional to the luminance registration of a corresponding region of the camera's photoreceptor. When congenitally blind subjects control the camera, they report having experiences as of three-dimensional shapes at a distance within as little as an hour. Subjects rapidly begin responding to parallax, occlusion, and looming (Bach-y-Rita, 2002). So, does TVSS show us that experiences of vibration/electrical stimulation against the skin can represent either vibrations/electrical stimulation or three-dimensional shapes in distal space? The answer seems to be no. Subjects report having *novel experiences* in addition to already familiar sorts of experiences of vibration/electrical stimulation of the skin. The two kinds of experiences compete for attention: When subjects attend to the stimulation of their skin, their abilities to respond to distal space suffer. Moreover, subjects' experiences of distal space remain constant even when the array is moved from their back to their stomach (Ibid.). This all strongly suggests that the subjects' novel experiences represent distal space and their already familiar experiences of vibration/electrical stimulation continue to play the same representational role that they always have.

constrain its possible contents. This is the causal covariation theory: Rs represent Os if Os typically cause Rs under normal circumstances. If this is how experiential representation works, then my conclusions will not follow from my arguments. Consider visual-shape displacement. If the causal covariation theory is true for experiential representation, then the kinds of experiences that we have in response to spheres will represent cubes when had by our visual-shape displaced twins, since cubes will be the normal cause of that sort of experience in them.¹⁹

Causal covariation accounts have their place. They are useful in understanding the sense in which individual cells in early vision can be said to represent oriented edges: Each one is selectively responsive to edges at a particular angle and thereby functions as an edge-at-that-angle detector. My arguments, however, concern representation in personal-level perceptual experience, not subpersonal neuronal mechanisms. And it is unlikely that perceptual experiences represent in the way that simple property detectors do; when applied to personal-level perceptual experience, the causal covariation account carries some extreme consequences.

David Chalmers employs a causal covariation account of perceptual representation, which he calls *the Fregean theory*, to argue that the brain-in-a-vat scenario is not a skeptical hypothesis:

[T]he perceptual experiences of a subject in a matrix [i.e., a brain-in-a-vat] are largely veridical. This picture coheres well with the Fregean picture ... according to which perception represents properties that are the normal causes of certain sorts of perceptual experiences. This Fregean content does not require properties of a highly constrained sort, so if it turns out that the relevant experiences are typically caused by certain computational properties, then those properties are represented by our experiences. (Chalmers, 2010, p. xxv).

As Chalmers notes, his striking conclusion follows from the minimalist constraints that Fregean/causally covariant content places upon what properties it can represent.

Most philosophers will be inclined to take Chalmers' argument as a reason to reject the causal covariation theory of experiential representation. Indeed, Terence Horgan, John Tienson, and George Graham all argue against a causal covariation account of experiential representation on the grounds that it cannot accommodate our intuitions that the brain-in-a-vat scenario is a skeptical scenario (Horgan, Tienson, & Graham, 2004; see also, McGinn, 1991). More generally, Tyler Burge argues against the causal covariation account of perceptual content on the grounds that it cannot allow for the possibility that certain kinds of perceptual states may typically (or always) misrepresent the subject's environment (Burge, 2010, Chapter 8; see also McGinn, 1989), such as experiences of a shade of green that is never instantiated. Chalmers' argument can be understood as an illustration of Burge's concern—if experiential content is causally covariant and we are brains-in-vats, then we do not systematically misperceive our environment. And finally, in his arguments for separatism, Papineau (2014) claims that a brain-in-a-vat in empty space will have conscious experiences that lack representational contents because it “doesn't have any environment to speak of” (p. 24). This suggests that Papineau

¹⁹Another example: Birds can see UV light; humans cannot. When birds look at UV gradients, they will have heterogeneous color experiences, humans will have homogenous color experiences, yet the causes of both sorts of experiences are identical. If experiential contents are causally covariant, then the contents of both experiences may be identical—this will depend upon how often UV gradients cause homogenous color experiences in humans. Thanks to an anonymous referee for offering this example and pressing this objection.

would reject Chalmers's claim that we can treat the brain-in-a-vat's computer simulation as a kind of environment that perceptual experiences can connect to in whatever ways are required for representation. But Papineau does not deny that the simulation causes the brain-in-a-vat's experiences. So, while Papineau does not presuppose any particular theory of representation in his arguments for separatism, he does presuppose that experiences do not always represent whatever typically happens to cause them. While my arguments require that we reject a causal covariation account of experiential representation, this is not a presupposition that my interlocutors are likely to contest, since embracing the causal covariation view of experiential representation carries some extreme consequences.²⁰

4 | METHODS FOR EXPLORING PHENOMENAL STRUCTURE

Phenomenal structure has received relatively little systematic exploration in the literature. Existing efforts employ an *outside-in* methodology: They utilize the results of psychophysics, neurophysiology, and multi-dimensional scaling to backwards engineer the structure of our experiences. C. L. Hardin pursues this sort of strategy for color perception (1988), and Austen Clark (1993) does so for perceptual experiences more generally. In Section 2, I employed an *inside-out* methodology, which relies upon the analysis of our intuitive reactions to thought experiments. Here I discuss how the two strategies relate.

The outside-in methodology promises to reveal structural features of our experiences that are too subtle to be noticed in the course of everyday life or through concerted introspection. For example, it turns out that human beings can discern more shades of green than red (Hardin, 1988). Facts like this are not part of our commonsense understanding of color experience and require careful experimentation to discover.

The inside-out method will tend to identify structural features of our experiences that we can recognize in our own experiences once they are pointed out. If we were not sensitive to structural features of our experiences in some way, then they could not drive our intuitive reactions to displacement and partial zombie cases. And the systematicity of our intuitions suggests that we have a tacit grasp of at least some of the structure of our experiences. There is, however, no a priori reason to think that this tacit understanding is complete. Outside-in methods promise to complement inside-out methods by locating phenomenal structures that evade our intuitive grasp.

5 | ARGUMENTS FOR SEPARATISM

Separatism and phenomenal schematics disagree over the existence of a priori connections between the phenomenality and intentionality of our experiences. To better understand the nature of this dispute and the more general lesson that phenomenal schematics has to offer, let us consider Papineau's central metaphor for understanding separatism: The comparison of phenomenal characters with the typographical properties of written words²¹:

²⁰There may be other theories of experiential representation that would impose minimalist constraints on the possible contents of experience. But details of such theories would need to be articulated before we could determine how they would rule on the cases discussed in Section 2.

²¹See Papineau (2014, pp. 1–2, 18, 2016, pp. 343, 339–340).

Written sentences are the outputs of processes designed to produce representations that will convey information to readers. It does not follow that all the properties of sentences are essentially representational. Their typographical properties are not ... It is entirely contingent that this arrangement of marks on paper means what it does ... I think the same about the relation between the conscious and representational properties of sensory experiences: The former stand to the latter just as the typographical properties of sentences stand to their representational contents. It is not essential to a given conscious experience that it stand for the truth condition it does. (Papineau, 2016, p. 332).

Papineau's comparison elegantly illustrates his view. While one must use some typography or other in order to represent using a written language, that does not mean that the typographical properties of written words are themselves representational properties, or that the typographical and representational properties of a particular word are any more than contingently related. Likewise, separatists claim, in order to have an experience that can represent whatever it does, that experience must have some phenomenal character or other, but that does not mean that the specific phenomenal character of that experience is any more than contingently related to that experience's representational content.

Papineau's metaphor brings to the fore the defining claim of separatism as it has been understood and defended in the literature:

Contingency: Intentionality and phenomenality are not merely distinct, they are contingently related.²²

As Papineau (2016) says, phenomenal characters, "have no constitutive connection with the ordinary objectual properties they contingently represent" (p. 344).

Here I review the three arguments that Kim and Papineau provide for separatism in order to show that these arguments cannot establish Contingency; at best they can only establish a weaker claim:

Distinctness: Intentionality and phenomenality are not identical to one another and neither is wholly grounded in or reducible to the other.

Notice that phenomenal schematics is compatible with Distinctness (but not Contingency). Since Kim's and Papineau's arguments can only establish Distinctness, those arguments do not decide between separatism and phenomenal schematics.²³

²²William Lycan (2008) describes separatism as the view that intentionality and phenomenality, "are unrelated, except possibly that some mental states happen to have both" (p. 240). Curtis Brown describes separatism as the view that, "any given phenomenal character could be accompanied by any intentional properties (or none), and vice versa" (2016, Section 3.4). And David Bourget and Angela Mendelovici characterize it as the view that "consciousness and intentionality do not bear interesting metaphysical relations to each other" (2017, Section 3.2).

²³Papineau backs away from Contingency at times, claiming that he is "happy to allow that the intrinsic nature of certain conscious sensory experiences suits them excellently for certain representational purposes" (2014, p. 27). And his reasons partly overlap with the arguments of Section 2.1.1. But Papineau's concessions are difficult to reconcile with his canonical statements of separatism, his repeated claims that phenomenality and intentionality are contingently related, and his characterization of phenomenal characters as vehicular properties of experiential representation, comparable to the color of a printed word.

Kim argues that we can in principle offer adequate functional analyses of intentional contents in terms of sensation, behavior, and cognitive processing; phenomenal characters, however, admit of no such functional analysis (2005, Chapter 6). According to Kim (2005), “qualia inversion is metaphysically possible” (Chapter 6, fn. 18). Hence:

[T]wo perceivers who behave identically with respect to input applied to their sensory receptors can have different sensory experiences ... qualia are not functionalizable, and hence physically irreducible. (Kim, 2005, p. 170).

If intentional contents can be naturalized via functional analysis and phenomenal characters cannot, then the two must be distinct.

To see why Kim's *naturalization asymmetry* cannot establish Contingency, notice that, as Kim himself claims, phenomenal structure may be functionally analyzable:

Intrinsic qualities of qualia are not functionalizable and therefore are irreducible, and hence causally impotent ... In contrast, certain important relational facts about qualia, in particular, their similarities and differences, are detectable and functionalizable. (Kim, 2005, p. 173).

If intentional contents and phenomenal structures can be functionally analyzed, then it is possible that at least some of the functions that our intentional contents reduce to are themselves partly composed out of the functions that phenomenal structures reduce to. If so, then phenomenal structure would constrain intentional content. Hence, Kim's argument can establish Distinctness, but not Contingency.²⁴

David Papineau (2014, 2016) offers two arguments for separatism. First, he appeals to a number of thought experiments and real-life examples to argue (a) that the phenomenal characters of our experiences are narrow (they supervene with metaphysical necessity upon the subject's physical/functional constitution) and (b) that the intentional contents of our experiences are wide (they supervene with metaphysical necessity upon both environmental and subject-internal factors). If intentionality is wide and phenomenality is narrow, then they must be distinct. Call this *the supervenience asymmetry argument*.

To claim that two phenomena are contingently related because one is wide and the other is narrow is not a valid inference. Larry's being the tallest man alive—a wide property—partly depends upon Larry's height—a narrow property. The two properties are not contingently related. Or, to borrow a counterexample from philosophy of mind, consider what content externalists can say about the relationship between the syntactic and semantic properties of a subject's mental states. Content externalists claim that both subject-internal and environmental factors play a role in determining the representational contents of a subject's mental states with metaphysical necessity. One way of understanding the role that subject-internal factors play (alongside environmental factors) in determining mental content is to claim that *the syntactic properties* of a subject's mental states are determined by the local, neurocomputational properties of the vehicles of that subject's mental states. In turn, the semantic properties of a subject's mental states will be determined by the ways in which the narrow, syntactically individuated

²⁴I suspect that Kim would be happy with this result. While Kim is the most cited example of a separatist in the literature (e.g., Bourget & Mendelovici, 2017; Brown, 2016; Horgan & Tienson, 2002; Lycan, 2008), his primary goal is to establish the naturalization asymmetry, not Contingency.

vehicles of the subject's mental states are causally regulated by the external environment. Because syntax can constrain semantics—a transitive verb cannot mean *Larry*—the narrow syntactic properties of the subject's mental states can constrain the wide contents of those mental states. For example, the concept WATER can refer to H₂O or XYZ depending upon how tokens of WATER are causally regulated by the environment. But WATER cannot refer to a person or be a logical operator, such as conjunction, because WATER has the functional-cum-syntactic profile of a natural kind concept (or a predicate, more generally) and not the functional-cum-syntactic profile of a name or a logical operator. So, on at least one way of understanding content externalism, syntactic properties and semantic properties are distinct because they take different supervenience bases, but narrow syntactic properties constrain wide semantic properties, nonetheless. The general form of inference in the supervenience asymmetry argument is invalid.

Papineau's (2014) second argument for separatism is his *abstraction asymmetry argument*: Any account of intentional content will have to understand intentional content as necessarily involving relationships to either propositions or properties understood as abstract objects, and, “there seems something quite amiss with the suggestion that my here-and-now conscious feelings are constituted by my bearing any kind of relation to abstract entities” (p. 6).²⁵ If intentional contents necessarily involve relationships to abstract objects, and phenomenal characters cannot, then they must be distinct.

Notice that a content externalist can grant that representational contents necessarily involve relations to abstract objects—though they are not obviously committed to this claim either—and still claim that the concrete here-and-now neurocomputational states of a subject constrain the representational contents of her mental states by virtue of determining the syntactic properties of the subject's mental states. Similarly, there is no reason why the phenomenal characters of our experiences could not constrain what abstract objects they can be related to. Papineau's abstraction asymmetry argument can only establish Distinctness.

Kim's and Papineau's arguments at best establish Distinctness, not Contingency. And Distinctness is compatible with phenomenal schematics.

6 | PHENOMENAL STRUCTURE AS THE GRAMMAR OF CONSCIOUSNESS

Papineau would have us believe that phenomenal characters are vehicular properties of experiential representation on par with the typographical properties of words. Yet the phenomenal structures of our experiences place formal and sometimes semantic constraints on the possible intentional contents that those experiences can have, and different structures impose different constraints (Section 2.3). In this respect, phenomenal structures are comparable to *the grammatical properties* of written words, not their typographical properties. Interpreted as a transitive verb, you cannot use “found” in a (grammatical) sentence and fail to represent an action that is performed by a grammatical subject upon a grammatical object. That is just what it is to be a transitive verb. Here we have both a formal constraint—any sentence that includes “found” must represent a grammatical subject and a grammatical object—and a semantic constraint—sentences including “found” must represent an action, as opposed to a mere happening, as in, “the battery died.”

²⁵Papineau also criticizes nominalist alternatives.

Similarly, our temperature experiences have the formal, scalar structure of a magnitude that acts upon the body surface or regions of the body volume. This structure imposes a formal constraint—such experiences are poorly suited to represent vectors, since they lack a second component that could serve to represent the vector's directionality—and also a semantic constraint—whatever property our temperature experiences represent must be capable of acting on the subject's body surface or body volume. If grammatical properties qualify as representational properties, then phenomenal structures do too.

There is, moreover, an important asymmetry between written words and phenomenal characters: Phenomenal characters possess their “grammatical properties” essentially. We can interpret “found” typographically—as an arrangement of colored shapes—in which case, “found” could have played the representational role played by any other English word. But we can also interpret “found” as a transitive verb, in which case “found” can be modeled as a function that only accepts ordered pairs of noun-phrases as arguments: It has an *argument structure* or *sub-categorization*, in the terminology of linguistics. Because we have a typographical way of thinking about words *in addition to* a grammatical way of thinking about words, it is conceivable that “found” might not have been a transitive verb; it could have been a noun, for instance.

What phenomenal schematics shows us is that the phenomenal characters of our experiences do not admit of anything like a typographical interpretation. We cannot conceive of alternative versions of our experiences that have argument structures different from the ones that they actually possess. Were we able to “think typographically” about the phenomenal characters of our experiences, then we should be able to imagine that what it is like for us to see cubes could have been what it is like for us to see spheres, and vice versa. After all, we have no difficulty in imagining that “sphere” could have meant cube and that “cube” could have meant sphere. But phenomenal displacement cases are only conceivable when the experiences involved have equivalent phenomenal structures (the inverted temperature spectrum, for instance). When it comes to thinking about the phenomenal characters of our experiences, the “grammatical interpretation” is *mandatory*.

What these observations suggest is that phenomenal characters are *intrinsically symbolic* in the sense that we can only conceive of them as having specific kinds of phenomenal structure that (a) determine what kind of role they can play in composing our total experiences at a time and (b) constrain the possible contents of our experiences. Papineau's claims that phenomenal characters, “have no constitutive tie to what they contingently represent” (2016, p. 343) and that “conscious sensory properties are not essentially representational at all” (2014, fn. 11) are false. An experience could not be the kind of experience that it is without having the specific phenomenal character that it does; we cannot help but conceive of phenomenal characters as having the specific phenomenal structures that they do; and the different kinds of phenomenal structure that our experiences can possess constrain the possible representational contents of those experiences in different ways. This is the central lesson of phenomenal schematics, and it marks a previously overlooked position that should be acceptable to all non-separatist philosophers of mind.

As noted in the introduction, phenomenal schematics remains neutral on how to answer:

- (2) Is there a kind of intentionality that suffices for the phenomenality of our experiences?
- (3) Does the phenomenality of our experiences suffice for a kind of intentionality?

Consequently, phenomenal schematics is compatible with both representationalism and phenomenal intentionalism, which respectively answer (2) and (3) in the positive. But phenomenal

schematics is not only compatible with these views, it promises to complement them by allowing them to offer more detailed specifications of the phenomenal-intentional relation.

Philosophers who answer (2) and (3) in the negative can also endorse phenomenal schematics. Call the resulting view *strong phenomenal schematics* (see Figure 1): While the phenomenal structures of our experiences place a priori constraints on the possible intentional contents of our experiences, the phenomenal characters of our experiences are not identical to, do not determine, and are not determined by the actual representational contents of our experiences.

Deciding whether to use phenomenal schematics to enrich representationalism or phenomenal intentionalism, on the one hand, or whether to opt for strong phenomenal schematics, on the other, is not a project that I shall pursue here. No matter which option one chooses, phenomenal schematics shows us that phenomenal characters mark a unique kind of representation, one whose structure is comparable to the grammar of a language or the rules of composition governing the representational elements used in diagrams, maps, or models. But phenomenal characters possess their grammar essentially. They are essentially symbolic.

ACKNOWLEDGEMENTS

Thanks to Terence Horgan, Shaun Nichols, Martin Davies, Jonathan Weinberg, Tim Bayne, Uriah Kriegel, Connie Rosati, Tyler Millhouse, Bryan Chambliss, Matt DeStefano, Caroline King, Robert Lazo, Max Kramer, Matt Schuller, and two anonymous referees from *Mind & Language* for providing some extremely helpful comments on this paper. I would also like to thank Jason Turner and Shaughan Lavine for helpful conversations.

ORCID

Brandon Ashby  <https://orcid.org/0000-0002-9897-9733>

REFERENCES

- Bach-y-Rita, P. (2002). Sensory substitution and qualia. In A. Noë & E. Thompson (Eds.), *Vision and mind* (pp. 497–514). Cambridge, MA: MIT Press. <https://doi.org/10.7551/mitpress/7111.003.002>
- Bayne, T. & Chalmers, D. (2003). What is the unity of consciousness? In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration, and dissociation* (pp. 23–58; reprinted in Chalmers, 2010). Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198508571.003.0002>
- Bourget, D., & Mendelovici, A. (2017). Phenomenal intentionality. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy* (Spring ed.). Retrieved from: <https://plato.stanford.edu/archives/spr2017/entries/phenomenal-intentionality/>
- Brown, C. (2016). Narrow mental content. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Summer ed.). Retrieved from: <https://plato.stanford.edu/archives/sum2016/entries/content-narrow/>
- Burge, T. (2010). *The origins of objectivity*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199581405.001.0001>
- Byrne, A. & Tye, M. (2006). Qualia ain't in the head. *Noûs*, 40(2), 241–255. <https://doi.org/10.1111/j.0029-4,624.2006.00608.x>
- Chalmers, D. (2006). Perception and the fall from eden. In T. S. Gendler & J. Hawthorne (Eds.), *Perceptual experience* (pp. 49–125; reprinted in Chalmers 2010). Oxford: Clarendon Press. <https://doi.org/10.1093/acprof:oso/9780199289769.003.0003>
- Chalmers, D. (2010). *The character of consciousness*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199289769.001.0001>
- Clark, A. (1993). *Sensory qualities*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198236801.001.0001>
- Dretske, F. (1996). Phenomenal externalism or if meanings ain't in the head, where are qualia? *Philosophical Issues*, 7, 143–158. <https://doi.org/10.2307/1522899>

- Farah, M. J. (Ed.) (2004). Visual form agnosia. In *Visual agnosia* (2nd ed., pp. 11–26). Cambridge, MA: MIT Press. <https://doi.org/10.7551/mitpress/7122.003.0003>
- Hardin, C. L. (1988). *Color for philosophers: Unweaving the rainbow* (expanded ed.). Indianapolis, IN: Hackett Publishing Company.
- Horgan, T. (2014). Phenomenal intentionality and secondary qualities. In B. Brogaard (Ed.), *Does perception have content?* (pp. 329–349). Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199756018.003.0014>
- Horgan, T. & Tienson, J. (2002). The intentionality of phenomenology and the phenomenology of intentionality. In D. Chalmers (Ed.), *Philosophy of mind: Classical and contemporary readings* (pp. 520–533). Oxford: Oxford University Press.
- Horgan, T., Tienson, J. & Graham, G. (2004). Phenomenal intentionality and the brain in a vat. In R. Schantz (Ed.), *The externalist challenge* (pp. 297–317). New York, NY: De Gruyter. <https://doi.org/10.1515/9783110915273.297>
- Kim, J. (2005). *Physicalism, or something near enough*. Princeton, NJ: Princeton University Press. <https://doi.org/10.1515/9781400840847>
- Kriegel, U. (2011). *The sources of intentionality*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199742974.001.0001>
- Loar, B. (2003). Phenomenal intentionality as the basis of mental content. In M. Hahn & B. Ramberg (Eds.), *Reflections and replies: Essays in honor of Tyler Burge* (pp. 229–258). Cambridge, MA: MIT Press.
- Lycan, W. (2008). Phenomenal intentionalities. *American Philosophical Quarterly*, 45(3), 233–252.
- Lycan, W. G. (2001). The case for phenomenal externalism. *Noûs*, 35(s15), 17–35. <https://doi.org/10.1111/0029-4624.35.s15.2>
- Masrour, F. (2014). The geometry of visual space and the nature of visual experience. *Philosophical Studies*, 172(7), 1813–1832. <https://doi.org/10.1007/s11098-014-0390-0>
- McGinn, C. (1989). *Mental content*. Oxford: Blackwell.
- McGinn, C. (1991). *The problem of consciousness*. Oxford: Blackwell.
- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83(4), 435–450. <https://doi.org/10.2307/2183914>
- Papineau, D. (2014). Sensory experience and representational properties. *Proceedings of the Aristotelian Society*, CXIV(1 Pt.), 1–30. <https://doi.org/10.1111/j.1467-9264.2014.00362.x>
- Papineau, D. (2016). Against representationalism (about conscious sensory experience). *International Journal of Philosophical Studies*, 24(3), 324–347. <https://doi.org/10.1080/09672559.2016.1175102>
- Searle, J. (1992). *The rediscovery of mind*. Cambridge, MA: MIT Press.
- Strawson, G. (Ed.) (2008). Real intentionality 3: Why intentionality entails consciousness. In *Real materialism and other essays* (pp. 281–306). Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199267422.003.0012>
- Tye, M. (2003). *Consciousness and persions: Unity and identity*. Cambridge, MA: MIT Press. <https://doi.org/10.7551/mitpress/2109.001.0001>
- Wagner, M. (2006). *The geometries of visual space*. London: Lawrence Erlbaum Associates. <https://doi.org/10.4324/9780203837627>

How to cite this article: Ashby B. Rainbow's end: The structure, character, and content of conscious experience. *Mind & Language*. 2020;1–19. <https://doi.org/10.1111/mila.12316>