

The limited role of particulars in phenomenal experience

by Neil Mehta

Forthcoming in the *Journal of Philosophy*. Please cite the final published version.

Consider two deeply appealing thoughts. The first is that we experience external particulars. I have seen Cirque du Soleil perform, and in doing so experienced *those particular acrobats*. I have listened to the impressionistic lyrics of Modest Mouse, and in doing so experienced *Isaac Brock's distinctive vocals*. I have drunk a bottle of cask-strength Macallan (not all at once!), and in doing so experienced *those very drams of whisky*.

The second thought is that what it's like to have an experience¹ – the **phenomenal character** of an experience – is somehow independent of external particulars. For what it's like to see Tweedledum is just the same as what it's like to see his identical twin Tweedledee. And what it's like to see Tweedledum is also just the same as what it's like to hallucinate a fellow – no fellow in particular – who looks just like Tweedledum.²

The first thought is readily captured by **phenomenal particularism**, the view that external particulars – such as particular external objects and external property-instantiations – are sometimes part of the phenomenal character of experience.³ The phenomenal particularist might hold that the phenomenal character of a given Macallan-drinking experience includes the particular dram of whisky itself, as well as particular property-instantiations such as its black cherry, raisin, and oak flavors.⁴ On this view, of course I experience external particulars: this occurs whenever those particulars are part of the phenomenal character of my experience.

The second thought is readily captured by **phenomenal generalism**, the view that external particulars are never part of phenomenal character. The phenomenal generalist might hold that the phenomenal character of a given Macallan-drinking experience includes only properties (not property-instantiations): perhaps properties of the subject's brain, properties of the whisky, or whatever. But the particular dram and its particular flavors are no part of

¹ 'What it's like' talk has been ubiquitous since Nagel (1974).

² I am non-committal about the nature of particulars – e.g., about whether haecceities underlie particularity. Admittedly, I will ultimately endorse claims like these: that particulars exist, that experiences of different particulars may have precisely the same phenomenal character, and that experiences of the same particular may have different phenomenal character. But such claims are compatible with a wide range of metaphysical theories of particulars. I thank an anonymous editor for prompting this clarification.

³ One might resist the idea that phenomenal character has parts. Still, something explains the similarities and differences in phenomenal character among an experience 'as of' a blue circle, an experience 'as of' a red circle, and an experience 'as of' a red triangle; my arguments will succeed on any reasonable characterization of the *explanans*. For convenience, I will continue to speak as though the *explanantia* are the parts of phenomenal character.

⁴ Phenomenal particularists include Brewer (2011), Campbell (2002), Fish (2009), Martin (2004) and (2006), and Sturgeon (2008).

phenomenal character.⁵ On this view, what it's like to have an experience is independent of external particulars because what composes phenomenal character is wholly non-particular. (For brevity, I will henceforth drop the adjective 'external' and speak simply of particulars.)

Since phenomenal particularism and phenomenal generalism are incompatible, we cannot adopt both. Since they exhaustively divide the possibilities, given that phenomenal character exists, we must adopt at least one. But it is hard to see how the phenomenal generalist can capture the first thought or how the phenomenal particularist can capture the second thought, and it is also hard to abandon either of these thoughts. There: that is the shape of the problem. I will argue that a novel version of phenomenal generalism provides a satisfying solution. Before doing so, let me introduce an assumption.

Traditionally, philosophers have held that all parts of phenomenal character are internal to the subject's mind – perhaps they are properties of the subject's soul, properties of the subject's brain, or mind-internal sense-data. Recently, many philosophers have instead held that almost all parts of phenomenal character are external to the subject's mind, if they are spatially located at all.⁶ On this view, when I see a ripe red mango, the phenomenal character of my visual experience might include the color, shape, and other seen features of that mango. I shall adopt without argument this externalist orientation towards phenomenal character.⁷ This does not unfairly tip the playing field, for virtually every philosopher who opposes phenomenal generalism accepts this orientation.

After introducing the phenomena which motivate phenomenal particularism (§1), this paper tackles both the negative task of revealing severe problems for phenomenal particularism (§2) and the positive task of outlining a satisfying phenomenal generalist alternative (§3). With respect to the negative task, the central charge will be that phenomenal particularism makes false predictions about similarities and differences in phenomenal character; specifically, that it mischaracterizes what it's like to have experiences that pick up wholly distinct features of what happens to be a single object. However, despite my phenomenal generalist loyalties, I am deeply sympathetic to the core motivations of phenomenal particularists. Consequently, with respect to the positive task, I aim not to rebut particularist motivations, but to capture them. Fully doing so will require careful reflection on the mental kinds under which experiences fall (§4-5). If my arguments succeed, then the phenomenal generalist can have her cake and eat *that particular* cake, too.

⁵ Phenomenal generalists include Chalmers (2006), Davies (1992), McGinn (1982), Pautz (2007) and (2009), Schellenberg (2010), and Tye (1995) and (2009).

⁶ The qualification 'almost' avoids fiddly problems, e.g., counterexamples involving headache experiences. (I thank Robert Schroer for bringing such cases to my attention.)

⁷ I argue for this orientation in Mehta (2013). My allies include Harman (1990), Martin (2002), and Tye (2002).

1. The particularity of perceptual experience

I am presently gazing out the window at a snow-covered tree and undergoing a corresponding visual experience. My experience of the snow-covered tree is particular in at least four ways.⁸

First, in virtue of undergoing this experience, I do not merely see *some tree or other* covered in snow. I see *that very snow-covered tree* and *its particular color, shape, etc.*⁹ Second, my experience of the snow-covered tree grounds a capacity for direct reference to that specific tree and its particular features: I can wonder how old *that tree* is, admire the *specific sparkling whiteness of the snow*, and so forth.¹⁰ Third, my experience grounds the capacity for direct perceptual knowledge about that tree and its particular properties – knowledge that the tree is snow-covered, that it has a certain shape and size, that its bark has certain characteristics, and the like.¹¹ Fourth, my experience introspectively seems ‘directed at’ this very tree and its particular features.¹² These humble claims about my experience are hard to deny.

On the simplest approach, these are not four disparate phenomena, but rather four symptoms of a single underlying phenomenon which we may call the **particularity of experience**. Happily, the phenomenal particularist can offer a simple explanation of this phenomenon: this particular tree, and the particular property-instantiations that I see, are *parts* of the phenomenal character of my experience. It would then be unsurprising that I thereby perceive those particulars; that my experience grounds the ability to directly refer to, and acquire knowledge of, those particulars; and that my experience introspectively seems directed at those particulars. By contrast, it is not immediately obvious how the phenomenal generalist may capture the particularity of experience so neatly – or at all.

Despite this seemingly powerful support, phenomenal particularism is untenable, as I shall next argue.

2. Against phenomenal particularism

Begin with the simplest version of phenomenal particularism, **pure phenomenal particularism**, according to which the phenomenal character of any experience is composed wholly of particulars – particular external objects and property-instantiations. At first glance, pure phenomenal particularism seems lovely: besides capturing the particularity of perceptual experience, it provides an austere and unified ontological approach to phenomenal character.

⁸ Sturgeon (2008) provides a helpful overview of these ways, though he divides them slightly differently.

⁹ See Tye (2009).

¹⁰ See Campbell (2002) and McDowell (1998). Notice that even if wondering about or admiring an item does not in general require the existence of that item, in these examples I still wonder about or admire a particular existing item.

¹¹ See McDowell (2008).

¹² See Martin (2006).

Unfortunately, there is a conclusive objection to this view: some imaginative experiences have phenomenal characters which clearly include no particulars. When I visually imagine a purple koala – but no koala in particular – there is something it's like for me to undergo the accompanying imaginative experience. But no particular koala or property-instantiation of purple is part of the phenomenal character of this experience. This is predictable, for this imaginative experience lacks all four symptoms of the particularity of experience: in virtue of undergoing this experience, I do not imagine any purple koala in particular; my experience does not ground direct reference to, or knowledge of, any particular purple koala; and my experience does not introspectively seem directed at any particular purple koala. Therefore, some phenomenal character must include non-particulars; indeed, some phenomenal character must be limited to non-particulars.

Let the phenomenal particularist retreat, then. According to **impure phenomenal particularism**, some phenomenal character includes particulars, and some phenomenal character includes non-particulars (plausibly, properties). Since successful perceptual experiences display particularity, the impure phenomenal particularist should hold that the phenomenal character of any successful perceptual experience includes particulars.¹³ But she should permit that the phenomenal character of some experiences, including some imaginative experiences, includes no particulars.¹⁴ The remainder of this section raises a *prima facie* and an *ultima facie* objection to impure phenomenal particularism; the latter also applies to pure phenomenal particularism.

The *prima facie* objection stems from the complexity of impure phenomenal particularism. As all disputants will concede, phenomenal character is a genuine, explanatorily significant mental kind. It is correspondingly costly to hold, as the impure phenomenal particularist does, that phenomenal character is ontologically disunified. Any simple view about the ontological nature of the parts of phenomenal character, such as the view that phenomenal character includes only properties, thus has a substantial *prima facie* advantage over impure phenomenal particularism.

Still, the impure phenomenal particularist might hope to minimize the complexity in her account. She may claim that the phenomenal character of any experience displaying particularity is composed wholly of particulars, while the phenomenal character of any experience lacking particularity is composed wholly of non-particulars.¹⁵ Although this account of phenomenal character remains disunified *across* experiences, it may at least secure ontological unity *within* any given experience: the phenomenal character of any given experience is composed wholly of either particulars or non-particulars.¹⁶ Further, the impure phenomenal particularist might claim, any disunity in the account of phenomenal character properly marks an important difference between particularity-involving and particularity-free experiences.

¹³ Even if a visual experience of the sky does not involve any particular objects, it still involves particular property-instantiations, like an instantiation of sky-blue.

¹⁴ Perhaps some imaginative experiences display particularity, as when I imagine my mother.

¹⁵ Fish (2009) defends a view in this vicinity.

¹⁶ Even this result may not be salvageable, for some experiences have both perceptual and imaginative elements.

I acknowledge the importance of this difference, which will figure prominently in my positive account, between particularity-involving and particularity-free experiences. Nevertheless, I will argue, the current view captures this difference only at the cost of neglecting potential similarities in phenomenal character across these two types of experiences.

Consider an *ideal imaginer* whose imaginative experiences are as vivid, stable, and forceful as her genuine perceptual experiences. Compare two experiences: her experience when she visually imagines a snow-covered tree (but no tree in particular), and her experience when she sees a snow-covered tree that happens to have just the features that she imagined. These experiences obviously have substantial similarities in phenomenal character. But the proposed version of impure phenomenal particularism cannot capture this, for the phenomenal character of the imaginative experience is composed wholly of non-particulars (like the non-particular properties of snowiness and whiteness), whereas the phenomenal character of the perceptual experience is composed wholly of particulars (like the particular snow and its particular instantiation of whiteness). The parts of these two phenomenal characters neither overlap nor fall under any single relevant type.

Admittedly, on the current proposal, some parts of the phenomenal character of the imaginative experience (e.g., whiteness) are instantiated by some parts of the phenomenal character of the perceptual experience (e.g., the instantiation of whiteness). But this does not suffice for genuine similarity in phenomenal character, for whiteness is not relevantly similar to an instantiation of whiteness.

Consequently, the phenomenal particularist should hold that the phenomenal character of particularity-involving experiences includes both particulars and properties, while the phenomenal character of particularity-free experiences is composed wholly of properties. Such a view can account for potential phenomenal similarities between experiences of these two types: these experiences may overlap in their non-particular parts. This view, however, remains ontologically messy.

This is a genuine objection, but, as I said, only a *prima facie* one. Let me proceed to the *ultima facie* objection, which applies to any form of phenomenal particularism. Broadly stated, the objection is that the phenomenal particularist makes implausible predictions about similarities and differences in phenomenal character. Before arriving at the most dialectically effective articulation of this objection, I must discuss some less dialectically effective ones.

One articulation of this objection focuses on experiences of distinct but perceptually indistinguishable objects.¹⁷ Suppose I see a fellow with a small red hat, a distinctively egg-shaped body, and a smart blue bowtie. As it happens, I am undergoing a perceptual experience of Tweedledum. A few minutes later, from precisely the same angle, I see Tweedledee, the identical (and identically dressed and posed) twin of Tweedledum, and undergo another perceptual experience. These experiences have precisely the same phenomenal character. But these phenomenal characters cannot have precisely the same parts, for each

¹⁷ This concern is expressed by Davies (1992, pp. 25-26), McGinn (1982, p. 48), and Tye (1995, p. 138).

experience involves wholly distinct particular objects and property-instantiations.¹⁸

However, the phenomenal particularist has leeway in interpreting truths about similarities and differences of phenomenal character: she may interpret these as truths about phenomenal character *types*, rather than phenomenal character *tokens*. My experience of Tweedledum, she will say, is composed of Tweedledum, Tweedledum's clothing, some particular property-instantiations, and perhaps some properties. While the phenomenal character of my experience of Tweedledee cannot have precisely the same *token* parts, it may yet have parts of just the same *type* – Tweedledee, Tweedledee's clothing, some particular property-instantiations, and perhaps some properties. This maneuver permits the phenomenal particularist to recover, under one salient interpretation, the claim that these experiences have just the same phenomenal character.

Another articulation of the objection compares perceptual and hallucinatory experiences. Suppose that I hallucinate a fellow with a small red hat, a distinctively egg-shaped body, and a smart blue bowtie. Intuitively, my hallucinatory experience has the same type of phenomenal character as my successful perceptual experience of Tweedledum. But the phenomenal character of this hallucinatory experience need not include any external particulars, for there need not be any external particulars of the right sort in my environment; and even if there are, they will not be appropriately related to my experience. Thus, if Tweedledum is part of the phenomenal character of my successful perceptual experience of Tweedledum, while the phenomenal character of my hallucinatory experience lacks any corresponding part, then these experiences cannot even have the same type of phenomenal character.

Phenomenal particularists typically embrace this result, however. They concede that the hallucinating subject may not be able to *distinguish* the phenomenal character of the hallucinatory experience from the phenomenal character of the successful perceptual experience, but deny that the hallucinatory experience *really has* the same type of phenomenal character as the successful perceptual experience. Indeed, some phenomenal particularists boldly assert that the hallucinatory experience has no phenomenal character whatsoever.¹⁹ While the viability of such positions is heatedly disputed,²⁰ I will soon argue that we can articulate the objection to phenomenal particularism without entering into the intricacies of those disputes. For now, note that this phenomenal particularist response gains traction only because hallucinatory experiences are defective, and the hallucinating subject is thus in a bad epistemic position. Perhaps this lends plausibility to the idea that the hallucinating subject may be deeply mistaken about the phenomenal character of her experience.

The stage is now set for the decisive articulation of the objection. This articulation focuses on pairs of completely successful perceptual experiences that

¹⁸ Crucially, the challenge is not to explain the *perceptual indistinguishability* of Tweedledum and Tweedledee. As Soteriou (2000, p. 185) notes, this challenge can be met by pointing out that my experience of each twin makes me aware of precisely the same external properties. Rather, the challenge is to explain *the sameness in phenomenal character of my experiences* of Tweedledum and Tweedledee.

¹⁹ See Fish (2009), Logue (2012), and Sturgeon (2008).

²⁰ Many who dispute the viability of such positions have taken inspiration from Johnston (2004).

have no phenomenal similarity but happen to feature the same particular object. Suppose that I see a glass of Bordeaux, undergoing a corresponding visual experience. Shortly afterwards, while I am looking away, my friend accidentally spills the wine into my lap, and I undergo a corresponding tactile experience. What it's like to undergo the visual experience is completely different from what it's like to undergo the tactile experience. For the phenomenal character of the visual experience pertains to the redness of the wine, the shape of the glass, and the like, while the phenomenal character of the tactile experience pertains to the temperature of the wine, the way it spreads over my lap, and the like.

The phenomenal particularist should hold that the same particular body of wine is part of the phenomenal character of both experiences. For these experiences display particularity: in virtue of undergoing each experience, I perceive the wine; I am in a position to directly refer to the wine, itself; I am in a position to know facts about the wine; and my experience introspectively seems directed at the wine, itself. So the phenomenal particularist predicts that there must be some similarity in phenomenal character between these experiences. But this beggars belief – surely the mere fact that the same wine happens to be involved in the visual experience and the tactile experience does not guarantee any similarity in what it's like to undergo them! To further dramatize the point, consider other experiences of that very wine that might have no phenomenal similarity to my visual experience: my auditory experience of the wine dribbling to the floor, my olfactory experience of the wine and its rich aromas, and so forth.

How might the phenomenal particularist respond? In response to cases of distinct but perceptually indistinguishable objects, the phenomenal particularist distinguished between tokens and types of phenomenal character. But in the present context, this only tightens the screws: since the case concerns a single token (and *a fortiori*, a single type) of wine, the impure phenomenal particularist must hold that the phenomenal character of my visual and tactile experiences overlap in token *and* type.

In response to cases of successful perceptual experiences and 'matching' hallucinatory experiences, meanwhile, the impure phenomenal particularist denied any alleged similarity in phenomenal character. While these experiences are introspectively indistinguishable, she suggested, they in fact differ in phenomenal character. This response drew its credibility from the fact that the hallucinating subject undergoes a defective experience and is consequently in a bad epistemic position. Because my visual and tactile experiences are paradigms of successful perceptual experiences, however, a parallel reply is not credible here.

Finally, the phenomenal particularist might hope to exploit the fact that the problematic case involves experiences in different sensory modalities. Might this explain why these experiences have no overlap in phenomenal character? It cannot. For similarly problematic cases can occur within a single sensory modality, as when I see the same person from the front and from the back.

The phenomenal particularist lacks further recourse. While we must account for the particularity of experience, we ought not do so by subsuming particulars into the phenomenal character of experience. Let us consider a fresh approach.

3. Simple phenomenal generalism

According to **simple phenomenal generalism**, the view defended in the remainder of this paper, phenomenal character is composed wholly of properties.²¹ As this section shows, simple phenomenal generalism, unlike phenomenal particularism, can do justice to both of the deeply appealing thoughts mentioned in the introduction while effortlessly clearing the problems raised so far.²²

It is easy to see that simple phenomenal generalism avoids the problems raised above for phenomenal particularism. I argued that pure phenomenal particularism cannot allow for the fact that an imaginative experience may have phenomenal character, even though its subject need not be appropriately related to any external particulars of the relevant sort. Such cases do not pose even a *prima facie* problem for the simple phenomenal generalist, who holds that phenomenal character is composed wholly of properties. For he may further hold that, in some imaginative experiences, uninstantiated properties are parts of phenomenal character. Additionally, I pointed out that impure phenomenal particularism is at a *prima facie* disadvantage in virtue of being an ontologically complex theory of phenomenal character. Simple phenomenal generalism, as an ontologically simple theory of phenomenal character, has no such disadvantage. Finally, I observed that phenomenal particularists falsely predict that two experiences must be phenomenally similar merely in virtue of sharing particular parts. Simple phenomenal generalism clearly makes no such prediction.

Nevertheless, simple phenomenal generalism seems poorly equipped to capture the particularity of perceptual experience. Reconsider my successful perceptual experience of Tweedledum. In virtue of undergoing this experience, I see Tweedledum and his features in particular. Further, this experience grounds a capacity for direct reference to Tweedledum and his features: I can wonder who *that* person is, admire *his style* in particular, and think that *he* dresses fashionably.²³ Better yet, this experience grounds the capacity for direct perceptual knowledge about Tweedledum and his particular features. Finally, when I introspect this experience, it seems to concern, not just some fellow or other, and not just some features or other, but *that particular fellow* and *his particular features*. The phenomenal particularist explains these truths by holding that Tweedledum and his particular features are part of the phenomenal character. But how can the simple phenomenal generalist explain them?

She might try to capture such truths by holding that, in perceptual experience, the properties which compose phenomenal character determine which particular(s) the subject is perceiving. Such an approach encounters well-

²¹ I in fact prefer the subtly different view, explored in Mehta (2012), that phenomenal character is composed wholly of **centered properties**, which, roughly speaking, are *de se* analogues of ordinary properties. However, I will ignore that wrinkle here, since the motivations for that view are orthogonal to the current debate, and the coming arguments will go the same way regardless.

²² Comparing simple phenomenal generalism to other versions of phenomenal generalism is both valuable and beyond the purview of this paper.

²³ This last thought would, of course, be false.

known difficulties in cases of severe misperception – say, if I in fact see a white cube to my left, but (because of trick lighting, distorting mirrors, and such) that cube looks like a red sphere directly ahead of me. Troublingly, in such cases the particular object which I see does not instantiate the properties, such as redness and sphericity, which presumably compose the phenomenal character of my experience. Alternatively, one might try to capture such truths via a causal theory of perception, but many complaints have been lodged against this theory.

I will skirt this conflict, supposing for the sake of argument that these approaches are unsuccessful. Even so, there is a simple explanation of the data: particulars may be parts of *experience*, even if they are not parts of the *phenomenal character of experience*. In other words, there may be *non-phenomenal* parts of experience. This rejoinder puts the phenomenal particularist and the simple phenomenal generalist on an even footing with respect to capturing the particularity of experience.

It is worth repeating this critical idea: we need not identify the particularity of *experience* with the particularity of *the phenomenal character of experience*. Or, putting the point another way, we need not identify an admittedly central part of experience – its phenomenal character – with the whole of experience.²⁴

We have now reconciled the thoughts that begin this paper. The first thought is that experience is particular: simple phenomenal generalism can capture this thought by holding that experience has particular parts. The second thought is that phenomenal character is independent of particulars: simple phenomenal generalism captures this thought by holding that phenomenal character has only non-particular parts.

This, however, exposes the simple phenomenal generalist to a *tu quoque* from the impure phenomenal particularist. For, I earlier complained, impure phenomenal particularism is an ontologically complex theory of phenomenal character. But simple phenomenal generalism has become an ontologically complex theory (not of phenomenal character but) of experience: some experiences have both properties and particulars as parts. How, then, can I claim the ontological advantage over impure phenomenal particularism?

I reply that the ontological complexity of simple phenomenal generalism carves nature at the joints. For all should agree that some experiences, such as certain imaginative or hallucinatory experiences, entirely lack particularity. When I imagine or hallucinate an apple, there need be no apple in particular that I experience. A fortiori, my experience need not ground a capacity for direct reference to, or knowledge of, ‘it.’ Further, while it may introspectively seem to me that my hallucinatory experience is directed at some particular apple, such introspective seemings are simply mistaken. Notably, when I imaginatively experience an apple – but no apple in particular – it will not introspectively seem to me that my experience is directed at a specific apple.

Thus, incorporating an earlier suggestion from the phenomenal particularist, I propose that there is a deep divide between experiences which have particularity and experiences which lack it. The simple phenomenal

²⁴ While I hold that all particular parts of experience are non-phenomenal, I leave open whether all non-phenomenal parts of experience are particulars.

generalist may mark this divide ontologically: particularity-involving experiences have particulars as parts, while particularity-free experiences have no particulars as parts. All the while, however, the simple phenomenal generalist retains an ontologically simple account of phenomenal character.

She further possesses a simple account of experience: something is an experience in virtue of having phenomenal character.²⁵ (Given my assumption of an externalist framework for phenomenal character, ‘having’ should not be glossed as ‘instantiating,’ but in some other way – perhaps as ‘providing acquaintance with.’) Consequently, while a *given* experience may be ontologically complex – it may have both particulars and non-particulars as parts – *what makes something an experience* remains ontologically simple.

The resulting account invites a version of the Many-Properties Objection.²⁶ This objection begins by noting that what it’s like to see a red square and a blue circle differs from what it’s like to see a blue square and a red circle. But the same properties figure in both experiences: redness, blueness, squareness, and circularity. So how can the simple phenomenal generalist capture the phenomenal differences between these experiences?

One reply is that two entities with the same parts may still differ in type if those parts are differently structured, as vividly illustrated by a watch and a heap of disassembled watch parts. Though the experiences just mentioned have the same properties as parts, those parts are differently structured. An alternative reply is that not all of the same properties figure in both experiences – e.g., the property *being red and square* is part of the phenomenal character of the first experience, but not the second.

4. Experience and mental kinds

This section considers a nagging suspicion about simple phenomenal generalism, the dispelling of which will reveal some underappreciated resources of this account. Put simply, the suspicion is that the simple phenomenal generalist misclassifies experiences into mental kinds. For mustn’t she concede that imaginative experiences and successful perceptual experiences may have just the same phenomenal character, and may thus belong to the same mental kind? Yet don’t such experiences plainly always belong to different mental kinds?

²⁵ Or, if you prefer, something is a *conscious* experience in virtue of having phenomenal character.

²⁶ Introduced by Jackson (1977).

The **Argument from Mental Kinds** embodies this suspicion:

- (1) Any two experiences with precisely the same phenomenal character belong to the same mental kind.
- (2) Any successful perceptual experience belongs to a different mental kind than any purely imaginative experience.
- (3) If simple phenomenal generalism is true, then there is a successful perceptual experience with exactly the same phenomenal character as a purely imaginative experience.

(Conclusion) Simple phenomenal generalism is false.

The first premise expresses a natural thought about mental kinds. Just as it is natural to classify beliefs fine-grainedly according to their specific contents, it is natural to classify experiences fine-grainedly according to their specific phenomenal characters. On this thought, the phenomenal character of any experience exclusively determines what **phenomenal kind** that experience falls under.

The second premise garners support from several directions. First, successful perceptual experiences have particularity in the fourfold sense discussed in §1 and thereby put the subject in direct contact with the external world. By contrast, imaginative experiences typically lack particularity and typically do not put the subject in direct contact with the external world. Moreover, successful perceptual experiences differ from imaginative experiences in functional role – e.g., in their characteristic relations to belief and action: compare the normal responses of a hungry person to seeing an apple or imagining one. Further, successful perceptual experiences differ in epistemic import from imaginative experiences: the former typically ground direct knowledge, or at least justified belief, regarding the external world, while the latter typically do not. Such differences mark a deep difference in mental kind, perhaps akin to the difference between believing and supposing.

To support the third premise, recall the ideal imaginer (§2), whose imaginative experiences are just as vivid, stable, and forceful as her successful perceptual experiences. Suppose she first sees a scene and then closes her eyes and vividly imagines a perfectly matching scene. On the simple phenomenal generalist approach, the phenomenal character of both experiences will apparently consist of just the same properties, structured in just the same ways. To be sure, the successful perceptual experience will have particulars among its parts, while the imaginative experience will not. But that does not undermine the claim about phenomenal character.

In reply to the Argument from Mental Kinds, the phenomenal generalist might resist premise (3). For while she must admit a substantial overlap between the phenomenal characters of a successful perceptual experience and a ‘matching’ imaginative experience, she may hold that successful perceptual experiences have a **presentational aspect** to their phenomenal character which imaginative

experiences lack. This presentational aspect might be either an additional part of phenomenal character or an aspect of how the properties that compose phenomenal character are structured.

However, the phenomenal generalist may wish to concede premise (3). For just as a belief and a desire may share precisely the same content while differing radically in their functional and evidential roles, so too, one might hold, may a successful perceptual experience and an imaginative experience share precisely the same phenomenal character while differing radically in their functional and evidential roles. I am sympathetic to premise (3) for just this reason. In any case, it will be a valuable exercise to see why the argument still fails.

Here is the vital observation: the Argument from Mental Kinds presupposes that whenever two experiences share precisely the same phenomenal character, each experience belongs to only one mental kind. For premise (1) refers to “the” mental kind to which such experiences belong. This presupposition, I will argue, is false. To be sure, because these experiences share precisely the same phenomenal character, they fall under the same *phenomenal* kind, just as the objector claims. Yet they may simultaneously differ with respect to *another* mental kind – a **teleological kind** determined solely by the **proper function** of the experience.

Let me first minimally characterize proper function via paradigmatic instances: roughly speaking, the heart has the proper function of pumping blood, the liver has the proper function of filtering toxins, and so forth. These are gross simplifications, to be sure, as the proper functions of the heart and liver are more specific and varied, but this rough gesture suffices for our purposes. Knowledge of such paradigmatic instances, perhaps supplemented by appropriate inferential dispositions, suffices to fix the reference of our concept of proper function, just as knowledge of paradigmatic instances of cats, perhaps supplemented by appropriate inferential dispositions, suffices to fix the reference of our concept of a cat. In these paradigmatic instances, an item’s²⁷ proper function stems from its evolutionary history. The present understanding of proper function is thus rooted in the biological sciences,²⁸ and the empirical usefulness of this concept is evidence that proper functions are genuine properties, carving the world at its joints. The proper function of an item thus solely determines what I will call its **teleological kind**: hearts and livers belong to different teleological kinds in virtue of having different proper functions.

Perceptual experiences and imaginative experiences differ in proper function and thus in teleological kind. For the former must put the subject in direct contact with the external world to fulfill its proper function, while the latter need not. This is evidenced by the fact that my visual experience “as of” a rattlesnake is successful only if that experience puts me in perceptual contact with an actual rattlesnake; otherwise, it is a failure, a merely hallucinatory experience. Not so for imaginative experience! The absence of nearby rattlesnakes in no way impugns my imaginative experience of a rattlesnake.²⁹

²⁷ Throughout, the term ‘item’ is used permissively. It refers to events, states, objects, or whatever.

²⁸ Here I follow Millikan (1984) and (1993).

²⁹ As I argue at length in Mehta (manuscript).

Here is the resulting picture. Consider a successful perceptual experience and a purely imaginative experience which share precisely the same phenomenal character. They belong to the same phenomenal kind, which explains their arresting subjective similarity. Yet in virtue of differing in proper function, they differ in teleological kind, which explains their differences in supporting knowledge, successful action, and the like. Unlike the view presupposed by premise (1), which classifies experiences along just one dimension, this nuanced view of mental kinds does justice to all of the relevant phenomena.³⁰

It is worth seeing how this approach interacts with the following assumption:

Fundamental Kind Assumption: Each mental item belongs to exactly one fundamental kind.

Many philosophers appear enamored of some assumption like this, as talk of an experience's 'fundamental kind' or 'sort' is commonplace.³¹ Surprisingly, explicit elaboration and defense of the Fundamental Kind Assumption has been thin. One suggestion, however, is that one kind is more fundamental than another whenever the former is more specific than the latter,³² and the most specific kind to which an item belongs is thus its fundamental kind.³³

As an aside, I would like to express skepticism about this defense of the Fundamental Kind Assumption. For notice that this motivation – that one kind is more fundamental than another whenever the former is more specific than the latter – does not entail the Fundamental Kind Assumption. After all, a single item may fall under several **cross-cutting kinds**, i.e., kinds such that none is more specific than the rest, as illustrated by the kind-demarcating properties *being mammalian* and *being able to swim*. The proponent of the Fundamental Kind Assumption must hold either (i) that there are no maximally specific cross-cutting kinds, or (ii) that whenever something falls under several maximally specific cross-cutting kinds, one of those kinds is more fundamental than the rest. The motivation for such views is unclear.

Notably, the existence of maximally specific cross-cutting mental kinds is a live philosophical possibility. For example, everyone agrees that some beliefs are not knowledge, and some also hold that not all knowledge requires belief – perhaps because knowing-how need not involve belief,³⁴ or because credences which fall short of full belief can be knowledge.³⁵ If knowledge does not require

³⁰ One might worry that teleological kinds are not mental kinds, for many non-mental items also have a teleological kind. So premise (2) – which claims that successful perceptual experiences and purely imaginative experiences differ in *mental* kind – is not yet vindicated. I reply that *some* teleological kinds – including the ones relevant here – are mental, since the proper functions in question are mental.

³¹ Examples include Crane (2006, p. 139), Martin (2006, p. 361), Neta (2008), and Snowdon (2005, p. 136).

³² This is eminently disputable, of course. Aristotle was a human male, and arguably human males comprise a genuine biological kind. But is this kind really more fundamental than the kind to which Aristotle belonged in virtue of being human, *simpliciter*?

³³ Martin (2006, p. 361) articulates this approach.

³⁴ See Brogaard (2012).

³⁵ See Moss (2013).

belief, then the mental kinds *believing that p* and *knowing that p* are good candidates for maximally specific cross-cutting mental kinds. Another example: everyone grants that some beliefs are not perceptions. Some also hold that some beliefs are perceptions: my seeing that there is a train ahead *just is* my believing that there is a train ahead. But some perceptions, one might further grant, are not beliefs: I might see something but disbelieve my eyes because I falsely believe that I am viewing an illusion.³⁶ On this view, *believing that p* and *perceiving that p* are good candidates for maximally specific, cross-cutting mental kinds. Now why suppose that, in each such case, one of these kinds must be more fundamental than the rest?³⁷

Still, on a conciliatory note, I observe that the Fundamental Kind Assumption is compatible with the simple phenomenal generalist view on offer. For though simple phenomenal generalism distinguishes phenomenal kinds from teleological kinds, it makes no claims about the fundamentality of any kind. Perhaps phenomenal kindhood is more fundamental than teleological kindhood, or vice versa. Or perhaps an experience's fundamental kind is a mental kind determined by *both* its phenomenal kind and its teleological kind.³⁸ I insist only that at least two mental kinds are in play.

In sum, because simple phenomenal generalism is minimally committal about how to classify experiences into kinds and non-committal about fundamental kinds, it can distinguish successful perceptual experiences from imaginative experiences via one mental kind, while still acknowledging that they may share just the same phenomenal character and thus share another mental kind.

5. What it's like to hallucinate

While the previous section examined the relationship between successful perceptual experiences and *imaginative* experiences, this section will consider the relationship between successful perceptual experiences and *hallucinatory* experiences. I will focus on the following claim:

Matching Thesis: For any successful perceptual experience E, there exists a purely hallucinatory experience with precisely the same phenomenal character as E.

The phenomenal particularist should deny the Matching Thesis, for successful perceptual experiences display particularity while (as argued in §3) purely hallucinatory experiences do not. To explain the particularity of experience, the phenomenal particularist appeals to particulars which compose phenomenal

³⁶ Lyons (2005, pp. 242-243) defends this view.

³⁷ Nanay (forthcoming) provides further reasons for skepticism about the Fundamental Kind Assumption.

³⁸ Advocates of this approach should exercise caution, for the intersection of two kinds need not itself be a kind.

character, so she should hold that particulars compose the phenomenal character of successful perceptual experiences but not of purely hallucinatory experiences.

Many, including myself, find the Matching Thesis very attractive, and even some phenomenal particularists acknowledge its allure.³⁹ But, the phenomenal particularist might think, one cannot accommodate the Matching Thesis while doing justice to the manifest particularity of successful perceptual experiences, so the Matching Thesis must be rejected. As a result, many phenomenal particularists have valiantly attempted to debunk the intuitive appeal of the Matching Thesis,⁴⁰ while opponents have fiercely resisted such attempts.⁴¹

Rather than joining the fray, I wish merely to observe that the simple phenomenal generalist can capture the particularity of successful perceptual experiences, as argued above, while simultaneously accommodating the Matching Thesis. The phenomenal particularist is barred from accepting the Matching Thesis because she holds that particulars are parts of the phenomenal character of successful perceptual experiences, but the purely hallucinating subject is not appropriately related to any particulars. However, the simple phenomenal generalist holds that the phenomenal character of successful perceptual experiences is composed wholly of properties structured in some way, and nothing bars her from holding that such properties structured in that way may compose the phenomenal character of a purely hallucinatory experience. Simple phenomenal generalism can therefore be readily paired with the Matching Thesis.

But this provokes a worry analogous to that discussed in the previous section: if, as the simple phenomenal generalist might propose, a successful perceptual experience and a purely hallucinatory experience may share precisely the same phenomenal character, then they will belong to the same mental kind. But any successful perceptual experience falls under a different mental kind than any purely hallucinatory experience. So the simple phenomenal generalist proposal is mistaken.

The simple phenomenal generalist cannot escape by holding that a successful perceptual experience and a 'matching' purely hallucinatory experience fall under the same phenomenal kind but differ in teleological kind. For successful perceptual experiences and purely hallucinatory experiences have the same proper function: to put the subject in direct contact with the world. To be sure, a purely hallucinatory experience does not *fulfill* this proper function, but it still *has* it.

The solution is to introduce another kind. To motivate this solution, consider a healthy heart and a diseased heart. While both share the same proper function and thus fall under the same teleological kind, they differ in fulfillment of this proper function. This difference is deep and explanatorily significant: it explains, e.g., why one heart circulates blood properly, but the other does not. In virtue of this difference, these two hearts differ in what I will call **appraisal kind**: the healthy heart is successful, while the diseased heart is unsuccessful.

³⁹ See Fish (2009, pp. 82-83).

⁴⁰ Representative examples include Fish (2009), Logue (2012), and Martin (2006).

⁴¹ Representative examples include Johnston (2004), Pautz (2007), Siegel (2010, ch. 6), and Tye (2009, pp. 112-113).

More generally, let us divide all items with a given proper function into those that fulfill that proper function and those that do not. This, I propose, divides those items into sub-kinds: the successful and the unsuccessful, relative to the given proper function. These are genuine kinds, for as illustrated by the heart example, there are typically deep explanatory differences between items that fulfill a given proper function and items which, while having the same proper function, do not. Appraisal kind thus stands to teleological kind as species to genus.

If this approach is correct, then a successful perceptual experience and a ‘matching’ purely hallucinatory experience, while falling under the same phenomenal and teleological kinds, differ in appraisal kind. For the former is successful, but the latter is not. This deep difference in kind grounds many significant differences between the successful perceptual experience and the purely hallucinatory experience. For example, the former, but not the latter, will (barring defeaters) ground *direct perceptual knowledge* and *non-accidentally successful action* with respect to the perceptible environment. Because appraisal kind is a subset of teleological kind, an appraisal kind will be mental whenever its corresponding teleological kind is mental. This account is non-committal about the relative fundamentality of phenomenal kinds, teleological kinds, and appraisal kinds.

What is the dialectical significance of the fact that simple phenomenal generalism, but not phenomenal particularism, is compatible with the Matching Thesis? Those sympathetic to the Matching Thesis should, of course, take this to favor simple phenomenal generalism. But for those who harbor active antipathy to the Matching Thesis, I underscore that simple phenomenal generalism merely *leaves room* for the Matching Thesis; it does not mandate it. Indeed, the mere fact that simple phenomenal generalism is less committal about the Matching Thesis than phenomenal particularism might itself be taken as a small advantage, for the simple phenomenal generalist does not stick her neck out unnecessarily.

6. Conclusion

The particularity of perceptual experience has driven many to phenomenal particularism. But, I have argued, simple phenomenal generalism can capture this particularity just as well as phenomenal particularism. Simple phenomenal generalism as developed here also has these advantages over phenomenal particularism:

- (a) It avoids the mistaken prediction that perceptual experiences overlap in phenomenal character merely in virtue of being experiences of the same particular object.
- (b) It provides an ontologically simple account of phenomenal character, not only for perceptual experiences, but also for imaginative and hallucinatory experiences: phenomenal character is composed wholly of properties.

- (c) It provides an ontologically simple account of what makes something an experience: something is an experience in virtue of having phenomenal character (where 'having' is not synonymous with 'instantiating').
- (d) It captures the deep similarities in kind among successful perceptual, imaginative, and hallucinatory experiences which arguably share the same phenomenal character, while still acknowledging that successful perceptual, imaginative, and hallucinatory experiences also differ deeply in kind.
- (e) It leaves room for, without mandating, the Matching Thesis: for any successful perceptual experience E, there exists a purely hallucinatory experience with precisely the same phenomenal character as E.

This is a powerful cumulative case for simple phenomenal generalism over phenomenal particularism.

ACKNOWLEDGMENTS:

I have accrued many debts in writing this paper. I am especially grateful for the generosity of Tamar Gendler, Pär Sundström, Todd Ganson, and Joseph Levine, who commented extensively on drafts. This paper also benefited from the helpful feedback of Andy Egan, Eric Lormand, Susanna Schellenberg, Andrew Bailey, Matt Walker, Robert Schroer, Damian Wassel, Cat Saint Croix, and participants in the 2012 Early Career Conference at the Northern Institute of Philosophy.

REFERENCES

- Brewer, B. (2011). *Perception and Its Objects*. Oxford: Oxford University Press.
- Brogaard, B. (2012). "Knowledge-how: a unified account." In *Knowing How: Essays on Knowledge, Mind, and Action*, eds. J. Bengson and M. Moffett (New York: Oxford University Press).
- Campbell, J. (2002). *Reference and Consciousness*. Oxford: Clarendon Press.
- Chalmers, D. (2006). "Perception and the fall from Eden." In *Perceptual Experience*, eds. T. Gendler & J. Hawthorne (Oxford University Press), pp. 49-125.
- Crane, T. (2006). "Is there a perceptual relation?" In *Perceptual Experience*, eds. T. Gendler and J. Hawthorne (Oxford University Press), pp. 126-146.
- Davies, M. (1992). "Perceptual content and local supervenience." *Proceedings of the Aristotelian Society* 66: 21-45.
- Fish, W. (2009). *Perception, Hallucination, and Illusion*. Oxford: Oxford University Press.
- Harman, G. (1990). "The intrinsic quality of experience." *Philosophical Perspectives* 4: 31-52.
- Jackson, F. (1977). *Perception: A Representative Theory*. Cambridge University Press.
- Johnston, M. (2004). "The obscure object of hallucination." *Philosophical Studies* 120: 113-83.
- Logue, H. (2012). "What should the naïve realist say about total hallucinations?" *Philosophical Perspectives* 26: 173-199.
- Lyons, J. (2005). "Perceptual belief and nonexperiential looks." *Philosophical Perspectives* 19: 237-56.
- Martin, M. (2002). "The transparency of experience." *Mind and Language* 17 (4): 376-425.
- Martin, M. (2004). "The limits of self-awareness." *Philosophical Studies* 120: 37-89.
- Martin, M. (2006). "On being alienated." In *Perceptual Experience*, eds. T. Gendler and J. Hawthorne (Oxford University Press), pp. 354-410.
- Mehta, N. (2012). "Exploring subjective representationalism." *Pacific Philosophical Quarterly* 93 (4): 570-594.
- Mehta, N. (2013). "Beyond transparency: the spatial argument for experiential externalism." *Philosophers' Imprint* 13 (8): 1-19.
- Mehta, N. (manuscript). "The normative role of perceptual experience."
- McDowell, J. (1998). "Having the world in view: Sellars, Kant, and intentionality." *Journal of Philosophy* 95: 431-491.
- McDowell, J. (2008). "The disjunctive conception of experience as material for a transcendental argument." *Disjunctivism: Perception, Action, Knowledge*, eds. A. Haddock and F. Macpherson (Oxford: Oxford University Press), pp. 376-389.
- McGinn, C. (1982). *The Character of Mind*. New York: Oxford University Press.
- Millikan, R. (1984). *Language, Thought and Other Biological Categories*. MIT Press.
- Millikan, R. (1993). *White Queen Psychology and Other Essays for Alice*. Cambridge, Massachusetts: MIT Press.
- Moss, S. (2013). "Epistemology formalized." *Philosophical Review* 122 (1): 1-43.

- Nagel, T. (1974). "What is it like to be a bat?" *Philosophical Review* 83 (4): 435-450.
- Nanay, B. (forthcoming.) "The representationalism versus relationalism debate: explanatory contextualism about perception." *European Journal of Philosophy*.
- Neta, R. (2008). "In defence of disjunctivism." In *Disjunctivism: Perception, Action, Knowledge*, eds. A. Haddock and F. Macpherson (Oxford: Oxford University Press), pp. 311-329.
- Pautz, A. (2007). "Intentionalism and perceptual presence." *Philosophical Perspectives* 21 (1): 495-541.
- Pautz, A. (2009). "What are the contents of experiences?" *Philosophical Quarterly* 59 (236): 483-507.
- Schellenberg, S. (2010). "The particularity and phenomenology of perceptual experience." *Philosophical Studies* 149: 19-48.
- Siegel, S. (2010). *The Contents of Visual Experience*. Oxford: Oxford University Press.
- Snowdon, P. (2005). "The formulation of disjunctivism: a response to Fish." *Proceedings of the Aristotelian Society* 105 (1): 129-141.
- Soteriou, M. (2000). "The particularity of visual perception." *European Journal of Philosophy* 8 (2): 173-189.
- Sturgeon, S. (2008). "Disjunctivism about visual experience." In *Disjunctivism: Perception, Action, Knowledge*, eds. A. Haddock and F. Macpherson (New York: Oxford University Press), pp. 112-143.
- Tye, M. (1995). *Ten Problems of Consciousness: a Representational Theory of the Phenomenal Mind*. Cambridge, Massachusetts: MIT Press.
- Tye, M. (2002). "Representationalism and the transparency of experience." *Noûs* 36 (1): 137-51.
- Tye, M. (2009). *Consciousness Revisited*. Cambridge, Massachusetts: MIT Press.