

6. Hallucination and sensation

The topics of this chapter are hallucination and sensation. I begin with an account of their natures (§1). After a brief interlude, (§2), I then lay out several arguments from hallucination: a basic argument from hallucination (§3), a screening-off argument (§4), and an argument from first-personal access (§5). I offer pluralist rebuttals of each.

1. Hallucination and sensation

The pluralist says that to be a perception is to be a sensory representation that represents by standing in a perceptual relation to one or more targets. The perception of a target is h(ard)-conscious if it is “lit up” with (i.e., partly constituted by) deep awareness; otherwise it is h-unconscious. But what might the pluralist say about hallucination and sensation?

The pluralist’s answer is inspired by the categorization account. She thinks that categorization is a sophisticated version of a much simpler capacity: the capacity to automatically and subpersonally *take* experiences in various ways – e.g., as a perception or a sensation. In simple creatures, taking a sensory representation as a perception might consist in nothing more being disposed to act as if the content of the representation obtains.

Drawing on this notion of *taking* an experience to be a particular way, the pluralist offers these accounts:

To be a *hallucination* is to be an involuntary sensory representation that is not a perception but is automatically and subpersonally taken to be one. In normal, mature human beings,

this occurs when and because the sensory representation is taken to have a perception-like gestalt.

To be a *sensation* is to be an involuntary sensory representation that is not a perception and is automatically and subpersonally taken *not* to be one. In normal, mature human beings, this occurs when and because the sensory representation is taken not to have a perception-like gestalt.

To give a high-altitude view of pluralism: perceptions, hallucinations, and sensations are alike in being involuntary. This separates them from sensory imaginings and episodic memories. What separates hallucinations and sensations from perceptions is that they do not involve a perceptual relation. (For that matter, what separates sensory imaginings from episodic memories is that they do not involve a memorial relation.) Finally, what distinguishes hallucinations from sensations is that the content of the sensory representation is taken to have a perception-like gestalt.¹

But I anticipate an objection. Suppose that you are about to give an important presentation. You are nervous and cannot help but vividly picture a contemptuous response from your audience. Because this experience occurs involuntarily, the pluralist must apparently count

¹ Note that what distinguishes hallucinations from sensations is not whether the experience *actually has* a perceptual gestalt. It is whether it is *subpersonally taken* to have one. For imagine an experience that does not have a perceptual gestalt but is mistakenly taken to have one – perhaps a phosphene experience that is taken to be a perception of mind-independent lights floating in the air. Such an experience is clearly a hallucination, not a sensation.

it as a hallucination or sensation. But the experience is actually a sensory imagining.² Or consider a war veteran who suffers from flashbacks: he periodically relives the experience of seeing a friend shot in the shoulder. Flashbacks are involuntary, so, again, it seems that the pluralist must count them as either hallucinations or sensations. But this is clearly incorrect: flashbacks are memories.

This is a *prima facie* compelling objection. But the pluralist has a powerful reply to it. She can say that she wants to understand the term *involuntary* in a particular way: as meaning *not having its content generated, in certain characteristic ways, by mental events outside the perceptual system*.³ This is a stipulation, not a claim about the meaning of an English word. The objection is then avoided, for it turns on experiences whose contents are generated by the subject's beliefs, desires, emotions, and so on. For instance, when a presenter imagines a contemptuous response from her audience, the content of her experience is generated by her beliefs, fears, and so on. Similarly, the contents of the flashbacks are generated by stored information about the events being remembered. (This information is within the perceptual system in a certain broad sense, but it is a narrower sense that I have in mind here.)

There is another concern. Consider an involuntary sensory representation that is not a perception: it might not be taken *either* as having *or* as lacking a perception-like gestalt. What should the pluralist make of it? I briefly suggest an answer in Appendix 6A.

The pluralist's theory of hallucination and sensation has many virtues. We have already seen one of them: the theory lets us explain why perceptions and matching hallucinations will, in the usual case, first-personally seem to present mind-independent particulars of the same kind.

² Beck (2018, p. 325) raises an objection like this. He uses it to suggest that some hallucinations that are generated purely internally, in the absence of any proximal stimulus, are better understood as sensory imaginings. It will shortly become clear why I disagree.

³ For an alternative response to an objection in this ballpark, see Ichikawa (2009, p. 107).

Again, the explanation is that a typical perception and a matching hallucination will involve deep awareness of precisely the same profile of sensory qualities. Thus both will be categorized as having a gestalt characteristic of presenting mind-independent particulars of the same kind. Relatedly, the pluralist's theory lets her give a precise characterization of what it is for a hallucination to match a perception: it is for both experiences to be sensory representations of exactly the same profile of sensory qualities. This theory also makes it easy to explain why an h-unconscious perception and a matching hallucination would cause any subject, whether simple or sophisticated, to form similar beliefs and take similar actions.⁴ This occurs because both experiences involve deep awareness of precisely the same sensory qualities. Consequently, they will have the same gestalts and be taken (and in the case of the sophisticated subject, categorized) in the same way.

Another virtue of the theory is that it allows for the possibility of h-unconscious hallucinations and sensations.⁵ After all, in a normal mature human being, an involuntary sensory representation will be taken as having a perception-like gestalt just on the basis of what sensory qualities are represented. This can occur whether or not the representation of these qualities is lit up with deep awareness – that is, whether or not the subject has partially revelatory access to those qualities.

The theory has one final virtue. It explains how it is possible for simple creatures to have hallucinations or sensations⁶: it is because they can *take* their experiences to be perception-like,

⁴ Holding fixed her background mental situation, of course.

⁵ Berger and Nanay (2016) argue that naïve realists have trouble explaining this datum. For some possible naïve realist responses, see Anaya and Clarke (2017).

⁶ It is *prima facie* difficult for some naïve realists to explain this datum. On this point, see Siegel (2004, §2).

even if their lack of conceptual sophistication makes it impossible for them to perform any categorization.

Conclusion. The pluralist's account of hallucination and sensation is theoretically powerful and intuitively plausible. It also has a surprising consequence. Consider the following views:

The *weak common kind view*: There is a joint-carving kind that includes all and only the perceptual experiences.

The *strong common kind view*: There is a joint-carving *fundamental* kind that includes all and only the perceptual experiences.

Representationalists often accept the strong common kind view. They almost universally accept the weak common kind view. However, the pluralist has strong reason to reject even the latter view: she cannot find a boundary worth drawing here.

To begin with, there might be a psychological kind formed by all and only the experiences that are genuinely presentational, but the pluralist (unlike, say, the sense-datum theorist) thinks that hallucinations do not present anything.⁷ More importantly, there might be a psychological kind formed by all and only the experiences that first-personally seem to be present mind-independent particulars. This is a standard representationalist reason for thinking that there is a joint-carving kind that consists of all and only the perceptual experiences. However, we have

⁷ Not even properties, since the relation of presentation must be perspectival and must update in real time. Only particulars can be presented.

seen that some perceptions, even though they in fact present mind-independent particulars, do not first-personally seem to do so. Thus they do not belong to this kind.

The pluralist does think that there is a joint-carving mental kind that includes *all* perceptual experiences: they are all instances of deep awareness. However, this kind does not include *only* perceptual experiences. It also includes sensations, episodic memories, and sensory imaginings. In this sense the pluralist accepts *some* sort of common kind view, albeit one that is far weaker than even the weak common kind view.

The weak common kind view is often thought to be very plausible. Is it an objection to pluralism that it rejects this view? Not at all. The standard motivations for the weak common kind view are that all perceptual experiences really present something (perhaps sense-data), or at least they first-personally seem to present mind-independent particulars. But we have strong theoretical reasons to reject the claim that all perceptual experiences really present something. And the pluralist offers direct counterexamples to the claim that all perceptual experiences first-personally seem to present mind-independent particulars: again, think of certain perceptions of faintly flashing lights or of fireflies in a grassy field. In this way, the pluralist undercuts the very motivations for the weak common kind view.

Now for other matters. Consider these views:

- (i) The *causal content view*, which says that any perceptual experience is a representation whose content concerns a causal relation between the target of the experience and the experience itself.
- (ii) The *perspectival content view*, which says that any visual perception represents its target's independence from the subject and its perspectival connectedness to the experience itself.

- (iii) The *dual component theory*, which says that to be a perception is to be a perceptual belief (or a perceptual seeming) that is caused by a sensation.

An important motivation for these views is that they can distinguish perceptions from sensations. In Appendix 6B, however, I argue that they draw the distinction in the wrong place. Pluralism does better. In Appendix 6C, I compare the pluralist's analysis of sensations to the view that sensations are perceptual experiences of a certain kind – for example, that afterimages are light illusions.

Finally, it is worthwhile to consider other kinds of experiences, such as sensory imaginings and episodic memories. What is the nature of these experiences? And why do they usually, though not always, first-personally seem to be the kinds of experiences that they are? I conjecture that these questions can be answered by extending the pluralist theory in some natural ways. See Appendix 6D.

2. Interlude

A theme of this book is that there is not just one problem of perception. There are many. Likewise, there are many arguments from hallucination. Still, such arguments usually begin in the same way: by asking us to compare two intrinsically identical subjects, only one of whom is in perceptual contact with her environment. For the sake of concreteness, let us suppose that the first subject enjoys a veridical perception of a tomato, call it the *tomato-perception*, which involves a certain neural event. The second subject suffers a matching hallucination, call it the *tomato-hallucination*, which involves a distinct but intrinsically identical neural event. The tomato-

hallucination is caused by the manipulations of a devious neuroscientist who ensures that its proximal cause is intrinsically identical to the proximal cause of the tomato-perception.

In the remainder of this chapter, I rebut three arguments against pluralism that focus on these subjects: a basic argument from hallucination (§3), a screening-off argument (§4), and an argument from first-personal access (§5). Along the way, I reveal some important advantages of pluralism over standard naïve realism.

3. A basic argument from hallucination

Consider:

The basic argument from hallucination:

- A1. The tomato-hallucination does not present the subject with a tomato.
 - A2. If the tomato-hallucination does not present the subject with a tomato, then the tomato-perception does not present the subject with a tomato.
-
- A3. The tomato-perception does not present the subject with a tomato.⁸

A1 is extremely plausible, since (we may stipulate) there is just no tomato for the hallucinating subject to be presented with. A3 is inconsistent with pluralism. Thus, the pluralist should reject A2.

⁸ Cite Robinson, Foster, etc.

You might try to support A2 by observing that that the neural events underlying the perception and the hallucination are effects of intrinsically identical proximal causes on intrinsically identical subjects. However, even assuming a certain picture of causation, this shows only that these neural events are intrinsically identical. It does not show that the tomato-perception must be intrinsically identical to the tomato-hallucination – but grant for the sake of argument that this is also the case. What is more crucial is that intrinsically identical causes may give rise to events that differ *in their relational aspects*. Thus the evidence is perfectly compatible with the possibility that the tomato-perception and the tomato-hallucination differ *in their relational aspects* – in particular, with the possibility that the tomato-perception presents a tomato and the tomato-hallucination does not.⁹ Lo and behold, this is the pluralist’s view.

But there is another line of support for A2: the tomato-hallucination first-personally seems to be the same as the tomato-perception in all qualitative respects, including relational ones. This does not *entail* that these experiences are qualitatively identical, since two entities can be radically qualitatively different even if they first-personally seem to be the same. But some kind of explanation is required here. The pluralist provides such an explanation: both experiences involve deep awareness of precisely the same sensory qualities. Given the subjects’ perceptual learning, both experiences then first-personally seem to be the same in certain relational aspects as well – both first-personally seem to present a tomato – even though only the perception has this relational aspect.

Still, this commits us to saying that first-personal appearances of presentation are sometimes misleading. So what, then, if the tomato-perception first-personally seems to present a tomato? Why think that this experience really does present a tomato? Here the pluralist will

⁹ On this point, see Langsam (1995, pp. 43-44); Martin (2004, pp. 55-58).

respond that the first-personal appearances can very reliable without being perfectly reliable. Her account of the epistemology of presentation, given above, is intended to fill out the details in a satisfying way.

All of this is reminiscent of how the naïve realist would respond to the argument.¹⁰ But there are two important advantages to the pluralist's response. The standard naïve realist view is that the tomato-perception and the tomato-hallucination might generate similar first-personal appearances, similar beliefs, and similar behaviors, but they are not similar in any *fundamental* respects. They are entirely different at the fundamental level, and a disjunctive explanation for the similarities is required.¹¹

By contrast, the pluralist honors *part* of the appearance that the tomato-perception and the tomato-hallucination are the same, indeed fundamentally so: she says that they provide exactly the same kind of deep awareness of exactly the same sensory qualities. Insofar as this is better than a pure error theory of the relevant first-personal appearances, this is an advantage for pluralism. What is more, the pluralist's categorization account provides a unified rather than disjunctive explanation of the similarities in the first-personal appearances, the doxastic effects, and the behaviors generated by the tomato-perception and the tomato-hallucination. This is a second advantage for pluralism over naïve realism: it is a theoretical virtue to provide simple, unified explanations rather than complicated disjunctive ones.

¹⁰ Again, see fn. 2.

¹¹ See Martin (2004, p. 43); Soteriou (2005, p. 178); Crane (2006, p. 139); Neta (2008, pp. 311-312); Fish (2009, pp. 94-95); Nudds (2009, p. 337); Brewer (2011, p. 94); Hellie (2013, p. 151); Logue (2013, p. 109); Genone (2016, p. 7); Gomes (2017, p. 534).

I conclude that the naïve realist and the pluralist alike can respond adequately to the basic argument from hallucination. But the pluralist's response is theoretically simpler, not to mention truer to the first-personal appearances.

4. A screening-off argument

Naïve realism is famously subject to this *screening-off argument*:

- B1. The tomato-perception and the tomato-hallucination will cause similar beliefs, actions, etc.
 - B2. If the tomato-perception and the tomato-hallucination will cause similar beliefs, actions, etc., then they belong to a common kind.
 - B3. If the tomato-perception and the tomato-hallucination belong to a common kind, then the *naïve kind* (the kind characterized by the naïve relation) does no explanatory work.
 - B4. If the naïve kind does no explanatory work, then naïve realism is false.
-
- B5. Naïve realism is false.¹²

¹² Naïve realist Michael Martin anticipates this argument against his own view (2004, pp. 61-62; 2006, pp. 369-372). Martin's preferred response is to reject B3 by accepting a deflationary view of what the common kind is (2004, §7; 2006, pp. 369-370).

Incidentally, Martin (2004, p. 62, fn. 17) attributes the argument to Sturgeon (1998) and (2000, ch. 1), but I believe that Martin is too modest; the argument is his own. What Sturgeon says is rather that if the naïve realist offers different explanations of apparent presentational character ("Scene-Immediacy") for perceptions and hallucinations, then she will "thereby [be offering a] disunifying theory" (1998, p. 198). This is evidently not a concern about screening-off.

What opens the naïve realist to this argument is that she takes the tomato-perception to belong to a proprietary kind – a kind that does not include the tomato-hallucination – where this kind has first-personal, doxastic, and behavioral significance. For the naïve realist, the kind in question is the *naïve kind*, the kind whose members are naïve relations.

The pluralist embraces a similar commitment. She also takes the tomato-perception to belong to a proprietary kind that has first-personal, doxastic, and behavioral significance. For the pluralist, this is what we might call the *perceptual kind*: the kind whose members are sensory representations that represent by standing in a perceptual relation to certain targets. Thus, the pluralist faces this parallel screening-off argument:

- C1. The tomato-perception and the tomato-hallucination will cause similar beliefs, actions, etc.
 - C2. If the tomato-perception and the tomato-hallucination will cause similar beliefs, actions, etc., then they belong to a common kind.
 - C3. If the tomato-perception and the tomato-hallucination belong to a common kind, then the perceptual kind does no explanatory work.
 - C4. If the perceptual kind does no explanatory work, then pluralism is false.
-
- C5. Pluralism is false.¹³

¹³ Relatedly, Benj Hellie discusses what he calls the *multimoderate view*: the view that the tomato-hallucination belongs to the fundamental kind defined by some common property *C*, while the tomato-perception belongs to the fundamental kind defined by the conjunctive property *being a visual perception of a tomato and being C*. Hellie takes the screening-off argument to be a good argument against this view (2013, p. 167). The multimoderate view is much like pluralism, although there is a subtle difference: the pluralist says that the tomato-perception belongs not to one conjunctive fundamental kind but to several distinct fundamental kinds. Still, the response that I will give on behalf of the pluralist should also save the multimoderate view.

Consider the premises of the argument, starting with:

- C1. The tomato-perception and the tomato-hallucination will cause similar beliefs, actions, etc.

The idea is to compare the effects of these experiences in two subjects who are otherwise as similar as possible. The claim is not that the relevant beliefs and actions will be *identical*: the tomato-perception might result in a genuinely singular belief about an actual tomato, whereas the tomato-hallucination will result only in a failed singular belief. That is why C1 says only that the effects will be *similar*. So understood, this premise is obviously true.

Next we have:

- C2. If the tomato-perception and the tomato-hallucination will cause similar beliefs, actions, etc., then they belong to a common kind.

The pluralist should grant this claim. For she accepts the antecedent, which is simply C1. And she accepts the consequent: she thinks that any perception and any matching hallucination belong to a common kind, indeed a *fundamental* one, insofar as they both involve deep awareness of exactly the same sensory qualities. She even thinks that the antecedent and consequent are connected: she holds that their membership in a common kind *explains why* the tomato-perception and the tomato-hallucination will have similar doxastic and behavioral effects.

We will get to C3 in a moment. First, however, skip ahead to:

C4. If the perceptual kind does no explanatory work, then pluralism is false.

This is indisputable. It is just built into pluralism that the perceptual kind helps to explain both the real presentational character and the strong singular character of typical perceptions. Thus, the only move available to the pluralist is to deny:

C3. If the tomato-perception and the tomato-hallucination belong to a common kind, then the perceptual kind does no explanatory work.

Fortunately, it is perfectly clear how the pluralist can justify rejecting this premise. Yes, these experiences share *some* of their first-personal appearances, doxastic effects, and behavioral effects. The pluralist grants that the common kind explains these similarities. But there are also important differences between the tomato-perception and the tomato-hallucination. The perception first-personally seems to present *that very tomato*; it positions the subject to form beliefs about *it*; and it positions the subject to act effectively on *it*. Nothing similar is true of the tomato-hallucination. These facts can be explained by appealing to a kind that is distinctive to the tomato-perception – say, by its being a sensory representation that represents by standing in a perceptual relation to that very tomato.

Some have argued that such a response only leads to further trouble. The concern is that we wish to account for the subject's *conscious perspective*, which is supposed to be constituted by what she actually perceives. But the hallucinating subject has a conscious perspective as well! So if the tomato-perception and the tomato-hallucination belong to a common kind, then

presumably it is this common kind that fully explains the subject's conscious perspective in both cases.¹⁴

I submit that the pluralist should reject the last step of this argument.

The pluralist grants that there is an important series of commonalities between the conscious perspectives involved in these experiences. The most fundamental commonality is that (i) both experiences involve deep awareness of exactly the same profile of sensory qualities. As a consequence, (ii) both experiences have a perceptual gestalt. Because of this, (iii) they will both first-personally seem to present various particulars, which will first-personally seem to be ordinary and mind-independent. This in turn brings it about that (iv) both subjects will be disposed to make similar judgments about their environments.¹⁵

All that being said, the pluralist does not think that the common kind *fully* explains the conscious perspective of the perceiver. She identifies an additional aspect of the perceiver's conscious perspective: the perceiver is presented with a particular tomato and its instances of redness. Presentation entails one kind of consciousness, *awareness consciousness*, which consists of *direct awareness-of*. A subject who is presented with the tomato and its instances of redness is thereby directly aware of them. For the pluralist, the perceptual kind is crucial to the explanation of *this* aspect of the perceiver's conscious perspective. There is no such aspect to the conscious perspective of the hallucinator.¹⁶

¹⁴ So says Michael Martin (2004, p. 64) and (2006, p. 371). He presses this concern against a naïve realist who rejects B3, but his concern would apply equally to a pluralist who denies C3.

¹⁵ Again, however, the hallucinator's judgments will be only putatively singular.

¹⁶ What happens if we interpret the screening-off argument as appealing to hard consciousness rather than awareness consciousness? Then the pluralist should reply that the tomato is not an aspect of consciousness in this sense. As I argued in chapters 3 and 5, this is independently plausible: there are no hard problems of consciousness associated with ordinary objects *per se*, nor with apparent presentation of them.

I conclude that the pluralist has nothing to fear from the screening-off argument.

5. An argument from first-personal access

Here is one last argument that threatens pluralism:

The argument from first-personal access:

- D1. The tomato-hallucination belongs to a positive, first-personally accessible mental kind (call it the *common kind*).
 - D2. For any positive mental kind *K*, if the tomato-hallucination belongs to *K*, then the tomato-perception belongs to *K*.
 - D3. If the tomato-perception belongs to the common kind, then its subject does not have first-personal access to the fact that her experience presents a tomato.
-
- D4. The subject of the tomato-perception does not have first-personal access to the fact that her experience presents a tomato.¹⁷

The argument is evidently valid. Since the pluralist rejects the conclusion of the argument, she owes a response.

The argument makes reference to *positive* mental kinds.¹⁸ The basic idea is that membership in a positive mental kind is earned by how an entity *is* rather than by how it *fails to*

¹⁷ This line of argument is pressed in Martin (2004, pp. 59-60).

¹⁸ My opponent might want to modify the argument to refer to mental kinds that are also *fundamental*. This would not affect my response.

be. As a foil, suppose that there is a mental kind defined by the property of not being first-personally discriminable from a veridical perception.¹⁹ This would not be a positive mental kind. It is difficult to characterize the notion of a positive mental kind more precisely than this, since it seems possible in principle to characterize anything either in terms of how it is or how it is not. But, in a concessive spirit, I will assume that this problem can be solved.²⁰

D1 says that the tomato-hallucination belongs to a positive, first-personally accessible mental kind, which we can label *the common kind*. This is the premise that naïve realists typically deny.²¹ But the pluralist embraces D1: she says that this hallucination belongs to a kind whose members are all instances of deep awareness of sensory qualities such as redness and roundness.

I regard this as a great advantage of pluralism over standard naïve realism, for I find D1 to be extremely plausible on both theoretical and first-personal grounds.

As far as theory goes, we want to explain why the tomato-hallucination will first-personally seem to present a mind-independent tomato, why it will first-personally seem to position the subject to refer to a tomato while knowing what she is thereby referring to, why it will position the subject to form a justified belief that “that thing” is a tomato, and why it will dispose a hungry subject to try to grab the apparent tomato. The simplest and most natural move is to appeal to a positive, first-personally accessible mental kind – preferably one that is shared by the tomato-perception, since the same phenomena will be found there. Other theoretical moves are also possible. But they are less simple and less natural.

¹⁹ See Martin (2006, p. 369).

²⁰ A *concessive* spirit, because the notion of a positive mental kind is only needed for an argument *against* pluralism. The pluralist has no particular use for the notion.

²¹ Cite Martin and others.

In addition, there is a first-personal motivation for D1: the tomato-hallucination will first-personally seem to belong to a positive mental kind, *even to the subject who knows that she is hallucinating*. We must provide an error theory of *some* of the first-personal appearances generated by this hallucination, and it is entirely possible to do that here, too. But it would be better to honor the appearance.²²

Now consider D2, which says that for any positive mental kind *K*, if the tomato-hallucination belongs to *K*, then the tomato-perception belongs to *K*. This is plausible because membership in a positive mental kind is earned by how an entity *is*, not by how it is not. But everything that the tomato-hallucination positively is, the tomato-perception positively is: we have stipulated that these experiences are intrinsically alike and have type-identical proximal causes. Perhaps the tomato-perception has additional positive aspects that are not found in the tomato-hallucination *H*, such as its relational aspects. But the reverse is surely not the case.²³

D1 and D2 entail that the tomato-perception belongs to the common kind. This brings us to the final premise of the argument, D3, which says that if the tomato-perception belongs to the common kind, then its subject does not have first-personal access to the fact that her experience presents a tomato.

²² Some naïve realists grant that hallucinations belong to a positive mental kind. For instance, William Alston conjectures that hallucinations are perceptual relations to mental images (1999, pp. 191-192). However, it is unclear how mental images could perceptually appear to be ordinary, mind-independent objects. More fundamentally, it is unclear what mental images are. (Are they sense-data? If so, Alston's theory incurs many familiar costs.)

A better suggestion is offered by naïve realist Keith Allen (2015), who argues that hallucinations are sensory imaginings that are mistaken for perceptions. I find this to be much more satisfactory than the standard naïve realist approach. But there is a problem. Sensory imaginings are voluntary, *in the technical sense described earlier in this chapter*: their contents are generated, in certain characteristic ways, by mental events outside the perceptual system. Hallucinations are not voluntary in this sense. Thus I prefer the pluralist's account.

²³ See Martin (2004, pp. 53-58) and (2006, p. 369).

The case for D3 is this. The subject of the tomato-perception will have first-personal access to the fact that her experience belongs to the common kind. After all, the subject of the tomato-*hallucination* is supposed to have first-personal access to the fact that *her* experience belongs to the common kind. Why would this access vanish when the subject is veridically perceiving? Suppose then that the subject of the tomato-perception also has first-personal access to the fact that her experience presents a tomato. Then she must have first-personal access to *two separate aspects* of her experience: the aspect associated with the common kind and the aspect associated with her being presented with a tomato. Evidently this is not the case.²⁴

The pluralist can block this argument. She agrees that there are two separable aspects to the tomato-perception. First, it is an instance of deep awareness of certain sensory qualities, such as redness. As a result, it belongs to a common kind with the tomato-hallucination. Second, the tomato-perception presents a tomato. The pluralist also agrees that each of these aspects is first-personally accessible to the subject.²⁵ She avoids the problem by adding that the subject's first-personal access to the fact that the tomato-perception presents a tomato is *routed through* her first-personal access to a perceptual gestalt generated by the common kind. In other words, the tomato-perception has two *metaphysically* separable aspects, but the subject's *epistemic* access to these aspects is not separable: her first-personal access to the fact that her experience presents a tomato epistemically depends on her first-personal access to the fact that her experience belongs to the common kind. The subject who suffers the tomato-hallucination still has first-personal access to the common kind. Why, then, does this not grant her first-personal access to the fact

²⁴ See Martin (2004, pp. 59-60).

²⁵ At least, given normal conditions. If the subject has misleading evidence that she is hallucinating, then she may lose her first-personal access to the fact that her experience contains a tomato.

that her experience presents a tomato? Because her experience does *not* present a tomato – you cannot have first-personal access to a “fact” that does not obtain.

In this way, the categorization account lets the pluralist evade the argument from first-personal access. The pluralist’s response also has an important advantage over standard naïve realist responses: the pluralist says that the tomato-hallucination belongs to a positive mental kind, indeed one that is first-personally accessible and shared with the tomato-perception. These claims are attractive on both theoretical and first-personal grounds.

6. Conclusion

It is well-known that hallucinations pose serious problems for naïve realism. There are no such problems for pluralism. Pluralism has the resources to block the basic argument from hallucination, the screening-off argument, and the argument from first-personal appearances. Perhaps standard versions of naïve realism can do the same. But pluralism can honor part of the appearance that the tomato-perception and the tomato-hallucination are fundamentally the same. It can give a simple, unified explanation of the phenomenological, doxastic, and behavioral similarities between them. And it can affirm that there is a first-personally accessible, positive mental kind that includes both experiences. The standard naïve realist cannot do any of these things.

Still, you might wonder whether standard representationalism honors our intuitions about hallucination even better than pluralism. The worry is that we do not merely find it plausible that any perception and any matching hallucination have *a* deep commonality. We find it plausible that they are *fundamentally the same*. Only the standard representationalist honors this bolder claim.

I agree that the bolder claim might be intuitive at first blush. But I say that, on reflection, the bolder intuition is overthrown. Perceptions play many important explanatory roles that hallucinations do not play. Perceptions really present us with trees and tables. They also position us to make singular reference to novel particulars in a way that positions us to know what we are thereby referring to. Hallucinations do not play these roles. Thus, on reflection it is plausible that there is a deep difference between any perception and any matching hallucination, in addition to the deep similarity. The pluralist honors the intuition of difference, as well.

Thus I say that pluralism fully honors *all* of our *reflectively stable* intuitions about perception, hallucination, and the relationship between the two.

The result is an account that is beautifully unified in many respects. There seems to be something that unites all sensory mental events, whether h-conscious or not: the pluralist explains this in terms of sensory representation. There seems to be something that unites all h-conscious events: the pluralist explains this in terms of deep awareness. There seems to be something tying together perceptions, illusions, hallucinations, and sensations: the pluralist says that they are all involuntary, in a particular technical sense. There seems to be something that unites all mental events with apparent presentational character, including typical perceptions typical hallucinations, and typical sensations – but some unusual imaginative experiences are included too, and some atypical perceptions are excluded. The pluralist explains this in terms of categorization via an overall gestalt. Yet there also seems to be something – real presentational character – that is distinctive to h-conscious perceptions: the pluralist says that this is a matter of deploying a sensory representation that represents via a perceptual relation and that is lit up with deep awareness. This package of views is elegant and explanatorily powerful.

From all of this, I conclude that pluralism is the best theory of perception, at least of those that have been examined in this book.

This conclusion is provisional: perhaps someone will find a better way to explain, or explain away, the many data that I have discussed. One option would be to experiment with a pluralistic variant of standard representationalism or standard naïve realism. I consider these approaches in Appendix 6E. There is also an important version of naïve realism that I have not yet discussed: *flexible naïve realism*, which holds that perceptions *and* hallucinations present real instances of sensory qualities. This is possible because sensory qualities are *ontologically flexible* – roughly, their existence can be fully grounded in either the mind or the world. I compare flexible naïve realism to pluralism in Appendix 6F. Finally, it is possible to construct a theory that is radically unlike those on the market. I would welcome such attempts.²⁶ I will stand by pluralism while I wait.

The next chapter mops up some remaining matters. In particular, it develops a pluralist explanation of why certain perceptions position us to form distinctive conceptions of objects and properties – for instance, why seeing a tomato positions us to conceive of redness and roundness as categorical and intrinsic.

Appendix 6A. Ambiguous experiences

What does the pluralist say about involuntary deep awareness that is not taken *either* as having *or* as lacking a perception-like gestalt? Suppose for example that I have sometimes seen faintly flashing lights, and I have also sometimes rubbed my eyes and just *seemed* to see faintly flashing

²⁶ For instance, Lisa Miracchi has developed a new *perception-first theory* (2017), and Justin D’Ambrosio has developed a novel version of adverbialism (2019). As these theories stand, I do not see how they would account for the data described in chapter 3. Still, it is possible to construct pluralistic versions of these theories, which I imagine would fare better. I leave it to the reader to consider these matters.

lights. The two experiences are first-personally indiscriminable. But then my brain is stimulated to produce faint images of flashing lights. Because of my perceptual learning due to my previous experiences, I do not take the experience either as presentational or as non-presentational. Call this an *ambiguous experience*. Is this a hallucination or a sensation?

According to the present account, it is neither. Such an experience is, so to speak, poised to tip over into hallucination or sensation, but until it does it belongs to a third category. Still, it is possible to accommodate those who think that ambiguous experiences are hallucinations: we can simply say that hallucinations are instances of involuntary deep awareness that are *not* taken to have a *non*-perception-like gestalt. It is also possible to accommodate those who think that ambiguous experiences are sensations. This time we would say that sensations are instances of involuntary deep awareness that are *not* taken to have a perception-like gestalt. But, for what it is worth, I prefer the account as written.

Appendix 6B. Arguments against some views about perception

Some of the data for pluralism that we have recently discussed simultaneously pose problems for three well-known views about perception.

First there is the *causal content view*. This view says that any perceptual experience is a representation whose content concerns a causal relation between the target of the experience and the experience itself. For example, when I see a blue vase, the content of my experience is not just <there is a blue vase in front of me>. It is rather something like <there is a blue vase in front of me that is causing this very experience>.²⁷ This content can be gussied up further if desired,

²⁷ This theory is introduced in Searle (1983, ch. 2).

for example by saying that the causation is automatic, linked to the present state of the object, and not mediated by any distinct experience.²⁸ In any case, a central motivation for positing a causal content is to distinguish perceptual experiences from mere sensations or sensory imaginings: these merely represent a target, without representing the relevant causal relation between the target and the experience.

This view runs into trouble. Consider again the case in which I rub my eyes and then open them; I experience both phosphenes and actual faint lights that are indiscriminable from phosphenes. The causal content theorist must say that my experience of phosphene-sensations differs in content from my perception of the actual faint lights. In particular, when I experience the actual lights, the content of my experience concerns a particular kind of causal relation between those lights and the experience itself; when I experience the phosphenes, the content of my experience does not concern such a causal relation. Surely this is not correct! If it were, I would be able to tell which targets were phosphenes and which were actual lights. I cannot do this – that is the crucial feature of the case.

Now consider a related view: the view that any visual perception represents the independence of the target from the subject and the perspectival connectedness of the target to the experience. Again, the idea is that when I see a blue vase, the content of my experience is not just that there is a blue vase in front of me. It is also part of the content that if I change my perspective on the vase, the vase will not *thereby* move. (The vase could move for other reasons, but not just in virtue of a change in my perspective.) In addition, it is part of the content that this change in perspective would, all else equal, cause a change in the phenomenology of my

²⁸ This is how the theory is developed in Millar (2014a, esp. p. 249) and (2014b, esp. pp. 642-643).

experience: I would experience whatever is on the other side of the vase. Call this the *perspectival content view*.²⁹

The perspectival content view shares a central motivation with the causal content view: a desire to explain the difference between perceptual experiences and mere sensations.³⁰ The thought is that when I have a sensation of a bright green after-image, I do not take the after-image to be something that I can get a different perspective on. That is what distinguishes it from a phenomenologically similar perceptual experience – say, a brief perception of a bright green light. The full force of the motivation can be brought out with an exotic example:

Suppose you are looking at a tiny doll. You take yourself to be in the usual sort of circumstance with respect to the doll... Moreover, you are correct: you are seeing a doll. You even play with the doll a bit, putting it into the little hands of its owner and then back to a shelf in front of you. Then your attention moves on to other things.

After an hour or so, however, something odd happens. You look back at the doll on the shelf and find that it seems to have lost its independence: it moves with movements of your head as if you were wearing a helmet with a imperceptible arm extending from the front, keeping the doll in your field of view. You hypothesize that someone has somehow attached the doll to your eyeglasses using a very thin string, without your knowing it.

... now suppose that the strange sequence of visual experiences continues in an even stranger vein. You decide to test the eyeglass hypothesis by moving your eyes without moving your head, and you find that the doll seems to move with your eyes as well. It seems to be sensitive to the slightest eye movement. And things get even stranger. When you close your eyes, you continue having a visual experience as of a doll. And when you try, with your eyes open, to put an opaque object right in front of the doll to block it from your view, your visual experience persists

²⁹ This view is defended at length in Siegel (2006).

³⁰ Ibid, p. 355.

in being a visual experience as of a doll. Overall, your experience of the doll comes to operate much like the experience of “seeing stars” from being hit on the head or from standing up too quickly. Just as nothing can occlude the “stars,” nothing can occlude the “doll”; and just as you can “see stars” while you are seeing other things, so too you continue to see things in the normal way even when the “doll” won’t leave your field of view. As with “seeing stars,” the apparent position of the “doll” is highly sensitive to eye movement.³¹

The story begins with *perceptual experiences* as of a doll and ends with mere doll-*sensations*. What makes the difference? According to the perspectival content view, it is that at the beginning of the story, you perceptually represent the doll as something that you could get a different perspective on. By the end of the story, you do not.

Unfortunately, the perspectival content view makes the same bad prediction as the causal content view. Return once more to the case in which I rub my eyes, open them, and experience both phosphenes and actual faint lights. The perspectival content view predicts that my experience of the phosphenes will have a different content than my experience of the actual lights: I will represent the lights, but not the phosphenes, as something that I could get a different perspective on. But then I ought to be able to tell which of the experienced targets are phosphenes and which are actual lights. I repeat that this is not possible.

Let us step back. There are deep similarities between the causal content view and the perspectival content view. Both are motivated by the explanatory need to distinguish perceptual experiences from mere sensations and by the observation that the subject is typically sensitive, in some sense, to certain sophisticated differences between the two. So far I am on board. Where I think these theories go wrong is in trying to analyze the subject’s sensitivity in terms of the

³¹ See Siegel (2006, pp. 369-370).

content of her experience. The pluralist instead analyzes this in terms of how the subject *takes* her experience to be.

The pluralist is thereby able to enjoy the advantages of these alternative theories. In the main body of this chapter, I explained in detail how it can capture the difference between perceptual experiences and sensations. Likewise, the pluralist can easily explain why, in the case described above, you shift from having perceptual experiences of a doll to having mere doll-sensations: it is because you cease to take your doll-experiences as having a perception-like gestalt. This occurs because the sensory qualities associated with your doll-experience do not interact with other sensory qualities in the usual ways. They do not change in the right ways as you move, block your view with opaque objects, or close your eyes. Instead, the sensory qualities associated with the doll persist no matter what – even while your experience has a perception-like gestalt with respect to other environmental objects.

The pluralist's account has a further advantage: it avoids a standard objection to the causal and perspectival content views. These views say that the content of perceptual experience is quite sophisticated. The problem is that the content of perceptual experience is supposed to be *conveyed to the subject* by the experience³² – but the relevant causal or perspectival features do not seem to be available to the subject. When I see a blue vase, I can attend to the vase and its blueness. But I cannot attend to the causal relation between the vase and my experience, nor can I attend to how changing my perspective on the vase will change my experience. This just does not seem to be part of what my experience conveys to me.

³² As Siegel (2010, pp. 42-44) rightly emphasizes, only what is conveyed to the subject is fit to be a *content*.

There are various replies available to those who posit sophisticated contents of experience.³³ But the pluralist avoids the problem altogether. She puts the sophistication into the process by which the subject's *subpersonal systems* take her experiences as being one way or another. Any such sophistication would of course be hidden from the subject, in the same way that facts about subpersonal syntactic processing are hidden from the subject.

These are my reasons for preferring pluralism to the causal and perspectival content views.

I will end with a brief mention of a third view: the *dual component theory*, which says that to be a perception is to be a perceptual belief or a perceptual seeming that is appropriately caused by a sensation. It should be immediately obvious that the dual component theory is subject to our focal counterexample: I might mistake an actual light for a phosphene, in which case I do not believe that there really is such a light. It does not even seem to me that there really is such a light. It is not that I have a belief or seeming that is somehow *masked* by other factors. It is that it would *require* other factors – in my view, some form of taking-as – for me to *acquire* such a belief or seeming in the first place.³⁴

Appendix 6C. Are sensations a kind of perceptual experience?

³³ See Siegel (2006, pp. 382-383); Millar (2014b, pp. 649-650).

³⁴ It is standard to reject the doxastic version of the dual-component theory on the basis of the problem of known illusion. The problem is that when we experience a known illusion, we apparently do not believe the illusory content of our experience, contrary to what the doxastic dual-component theorist predicts. Jake Quilty-Dunn offers an interesting response to this problem: effectively, his idea is that perceptual beliefs can be masked by the subject's knowledge (2015, §3.2-§3.3). I believe that the objection I have given circumvents this response.

The dominant view is that sensations are a novel kind of experience apart from perceptual experiences. The pluralist agrees. However, some have argued that sensations are just perceptual experiences of various kinds. For example, according to the *light illusion account*, afterimages are illusions (or perhaps hallucinations) in which we represent the existence of a mind-independent light that is projected from the subject's location.³⁵ The basic case for the light illusion account is that in many respects, the perceptual system processes afterimages as though they were mind-independent. For example, afterimages can appear to show a certain degree of independence from the movements of the subject's eyes, they can appear to be occluded under certain circumstances, and they can appear to be visible from different perspectives.³⁶ These points are not evident from the armchair, but they are established by careful empirical research, and it is possible to extend them to other sensations such as phosphenes.³⁷

I grant that this is strong evidence that the *perceptual system* treats afterimages as mind-independent. What I contest is whether it is strong evidence that the perceptual system represents afterimages *to us* as mind-independent. There is much that happens in subpersonal perceptual processing that is never made manifest to the subject. I believe that armchair reflection suffices to show that sensations *often* (but not always) do not strike us as bodies of light.³⁸ What is crucial to this response is the distinction between what is built into our sensory representations themselves and what we make of those representations at the personal level.

³⁵ This account is defended in Phillips (2013, p. 427 and p. 433).

³⁶ Ibid, §9-§11.

³⁷ Ibid, pp. 446-447.

³⁸ I do agree with the light illusion theorist that afterimages can sometimes be first-personally mistaken for perceptions of mind-independent targets: that was a central part of my own case for various aspects of the pluralist's account.

Appendix 6D. Some conjectures about other types of experiences

Here I offer some conjectures about how to extend pluralism to cover sensory imaginings and episodic memories. I begin with a pair of conjectures about the natures of these mental events:

To be an *episodic memory* is to be a voluntary sensory representation that represents one or more particulars by standing in a memorial relation to them – a relation that is mediated by stored perceptual information.

To be a *sensory imagining* is to be a voluntary sensory representation that is not an episodic memory.³⁹

I am using the term “information” factively. Otherwise the conjectures would be subject to obvious counterexamples.

An initial observation is that the kinds of sensory representations discussed so far – perceptions, hallucinations, and sensations – have all been involuntary. In contrast to this, sensory imaginings and episodic memories are *voluntary* sensory representations. Again, this term is being used with a stipulative meaning, not its ordinary English one: the claim is that these sensory representations have their contents generated, in certain characteristic ways, by mental states or events outside the perceptual system.

Another observation is that, typically, episodic memories first-personally strike us as episodic memories, and sensory imaginings first-personally strike us as sensory imaginings.

³⁹ I am open to other ways of specifying what kind of voluntary sensory representation is constitutive of sensory imagination.

What explains this? It is natural for the pluralist to respond by extending her remarks about taking-as. In particular, I offer these conjectures:

For an experience to first-personally seem to be an episodic memory or sensory imagining is for the subject to take the experience as such. In normal, mature human beings, this taking occurs on the basis of a distinctive *memorial* or *imaginative profile*.

The idea is that the subject automatically and subpersonally attempts to keep track of various facts about a sensory representation. For example: what judgments, intentions, and so on preceded the representation; what gestalts are associated with the sensory qualities represented, not only in the representation itself, but also in immediately prior sensory representations; and so on. On this basis, in any typical case, the subject automatically and subpersonally takes each sensory representation to be a certain kind of experience – a perception, sensation, episodic memory, sensory imagining, etc. A sophisticated subject will also normally categorize her experiences as belonging to one of these experiential kinds, using a concept that she develops as she matures.

I have already discussed what cues might indicate that a sensory representation is a perception or sensation. What might indicate that the representation is an episodic memory or sensory imagining, though? In both cases, *voluntariness* (in the technical sense described above) would be one important cue. The idea is not that voluntariness can be detected directly. The idea is rather that the content of the sensory representation can be compared with the content of recent beliefs, intentions, and so on, to arrive at a verdict about voluntariness. In addition, to distinguish sensory imaginings from episodic memories, the subject's subpersonal systems might

directly keep track of whether the content of the representation was generated by intentions to remember, thoughts about the past, etc.

These conjectures can explain many data about episodic memories and sensory imaginings.

One datum is that episodic memories are usually associated with a sense of pastness. The present account can explain what this consists in: it consists in taking the representation to be an episodic memory. It will be part of a sophisticated subject's theoretical understanding of what episodic memory is that its content concerns her own past.

Another datum is that it is possible for the subject to miscategorize her experiences. She might fabricate wholesale an episode from her childhood in which she is riding a merry-go-round, while mistaking this for an episodic memory. Alternatively, an actual memory of riding a merry-go-round might surface in the subject's mind, but she might dismiss it as a mere imagining. The present conjectures explain how this is possible. Miscategorizations can occur when the cues associated with an experience are misleading, when these cues do not give enough information one way or another, or when the cues are just processed incorrectly.

Relatedly, the present conjectures suggest a way of understanding what it is to be a memorial experience: it is to be *either* a memory *or* an experience that is subpersonally categorized as a memory. This suggestion is, of course, parallel to the pluralist's analysis of what a perceptual experience is. And, as with that analysis, the disjunction is not redundant: some real episodic memories do not first-personally seem to be episodic memories.

To develop an earlier example of this in more detail, consider a woman who is taking her children to ride a merry-go-round. As soon as she sees this scene, an image floats into her mind from the subterranean depths of her mind. It is an image of herself, as a small child, riding a merry-go-round, albeit one that looks quite different from the one she sees before her now. We

may suppose that her experience does not for a moment strike her as a memory. It seems to her that she is merely imagining things – that she never went on a merry-go-round herself, but is feeling nostalgia for the carefree days of childhood. In fact, however, she *did* go on a merry-go-round, one that looks just like what she is now picturing. More than that, the content of her current experience is generated by a stored perceptual representation formed when she rode that merry-go-round decades ago. In this case, she is having an actual episodic memory, but it does not even first-personally seem to her to be an episodic memory.

There are still other data to be explained. To begin with, the first-personal appearances are sometimes silent about whether an experience is an episodic memory or a sensory imagining – for example, when an image of a merry-go-round floats to the surface of the subject’s mind, she might be unsure whether she is remembering or imagining. In addition, the first-personal appearances might persist even when the subject knows that these appearances are mistaken. For instance, she might have a vivid false “memory” of what happened at her eighth birthday party. This might continue to seem like a genuine memory even after she sees a video of the party that depicts events that do not at all match her experience. Moreover, a typical memory and a matching false “memory” will have precisely the same mental effects. It is also possible to have unconscious memories and sensory imaginings. Finally, it is possible for conceptually unsophisticated creatures to have memories and sensory imaginings.

These data are perfectly parallel to data that I discussed regarding perception, hallucination, and sensation. The present conjectures allow for parallel explanations of these data.

Appendix 6E. Some other pluralistic theories

According to the pluralist, perception constitutively involves two very different kinds of awareness: non-revelatory awareness of particulars and partially revelatory awareness of universals. You might accept this much without accepting pluralism. For instance, you might accept *pluralistic representationalism*, which says that both of these kinds of awareness are representational, or *pluralistic naïve realism*, which says that both are non-representational. Here I consider these alternative theories.

Pluralistic representationalism. This theory would accept something like the pluralist's account of sensory representation. But it would say that deep awareness is representational, too. The defining characteristic of representations is that they are capable of getting things wrong, so the point of pluralistic representationalism would be to allow that we might misrepresent the natures of the sensory qualities, *understood as universals*.⁴⁰ However, I find it plausible that experience are always right about the natures of the sensory qualities. Thus I think that it is idle, even harmful, to allow for the possibility of misrepresentation.

The pluralistic representationalist might suggest that we *occasionally* misrepresent the very natures of the sensory qualities – for instance, when we experience supersaturated red, or when we suffer from the waterfall illusion and experience something as moving and not moving. I myself am not persuaded on these points. Perhaps it is nomologically impossible for anything to be supersaturated red. But simply on the basis of reflecting on my experiences, I find it very plausible that supersaturated red is a real property, one that it is instantiated in some metaphysically possible worlds. As for the waterfall illusion, one possible explanation is that we

⁴⁰ Does it support pluralistic representationalism that we sometimes get wrong, not the *natures* of the sensory qualities, but whether or not they are instantiated? I do not believe so. There is good reason to posit sensory representations of *instances* of the sensory qualities – we can surely be aware of such instances, even without h-consciousness – and this already makes it possible to explain how experience can be wrong about whether various sensory qualities are instantiated.

represent *distinct particulars* as moving and not-moving. Another possible explanation is that these experiences correctly reveal to us the natures of motion and stillness, which we then *incorrectly represent* as occurring in a single particular.

These are my reasons for preferring pluralism to pluralistic representationalism. I acknowledge that they are suggestive, not decisive.

Pluralistic naïve realism. This theory would accept the pluralist's account of deep awareness. But it would analyze perceptions (including h-unconscious ones) in terms of a non-representational *and non-revelatory* relation of direct awareness – call it *acquaintance* – to particulars. The core of the theory might be this:

To be a *perception* is to be a mental event that constitutively involves a relation of acquaintance to particulars.

To be a *hallucination* is to be a mental event that is not a perception but is mistakenly taken to be one.

For any mental state or event to be *h-conscious* is for it to constitutively involve deep awareness of one or more sensory qualities, understood as universals.

I believe that pluralistic naïve realism has many advantages over standard naïve realism. In particular, the pluralistic naïve realist can say that h-conscious hallucinations have a positive nature: they make the subject deeply aware of certain sensory qualities. This lets her honor the intuition that any perception and any matching hallucination are deeply similar, since she says that the perception makes the subject deeply aware of the very same sensory qualities. The

account also lets her explain why we mistake matching h-conscious hallucinations as perceptions: it is because both kinds of experiences have exactly the same gestalts – perception-like gestalts – and are thus taken in the same way. It also makes it easy for her to explain how simple creatures can hallucinate, for even simple creatures can *take* their experiences to be perceptions.

However, I do not see any advantages in cleaving to pluralistic naïve realism *rather than pluralism*.

The usual motivations for a naïve realist view are to account for real presentational character, apparent presentational character, and strong singular character. However, I have already argued that the pluralist explains these phenomena even more effectively than does the naïve realist. The pluralistic naïve realist might borrow these explanations, but of course that does not give her any comparative advantage.

At the same time, pluralistic naïve realism does have some comparative disadvantages. One is that it is plausible that there is some feature that unites all sensory events – perceptions, hallucinations, sensory imaginings, and so on – *whether they are h-conscious or not*. It is plausible that there is a single feature that makes all of these events *sensory*. The pluralist identifies such a feature: all sensory events constitutively involve sensory representations. The pluralistic naïve realist cannot say this, since she does not think that perceptions constitutively involve representations at all. She cannot say that all sensory events involve acquaintance with particulars: hallucinations and sensory imaginings do not. Nor can she say that all sensory events involve deep awareness: h-unconscious sensory events do not. It is not clear what other unified account she can give.

Pluralistic naïve realism has another comparative disadvantage: its account of h-unconscious hallucinations is unsatisfying. H-unconscious hallucinations are mistakenly taken to be perceptions, but why? What is it about them, positively speaking, that causes them to be taken

in this way? It is unclear what the pluralistic naïve realist can say in response. Even if she does provide an answer, this answer will be very different in kind from her answer for h-conscious hallucinations. *These* are taken as perceptions because they provide deep awareness that has a perception-like gestalt. Nothing like that can be said of h-unconscious hallucinations, which, according to the pluralistic naïve realist, do not involve any kind of awareness or representation of sensory qualities.

The pluralist, by contrast, provides a unified answer to the question. She says that in a normal, mature person in normal circumstances, *any* hallucination, whether h-conscious or h-unconscious, will be taken as a perception on the basis of the overall gestalt of the underlying sensory representations. This explanation is the same whether or not the hallucination is h-conscious. The only difference is that if it is h-conscious, then the sensory representations are lit up with deep awareness, thereby bringing certain aspects of the gestalt within the realm of hard consciousness.

A final problem arises for sensory imaginings. At the very least, it is metaphysically possible for these to occur without hard consciousness. What *are* h-unconscious sensory imaginings? They do not involve acquaintance with particulars, since they are imaginings, nor do they involve deep awareness, since they are h-unconscious. The natural move, then, is for the pluralistic naïve realist to say that they are *representations*, of a proprietary kind, of relations of acquaintance to particulars. The lesson is that the pluralistic naïve realist will likely want to posit a proprietary kind of representation, *in addition to* the relations of acquaintance and deep awareness. Pluralism is more parsimonious. It posits only sensory representation and deep awareness.

In sum, I see several considerations that favor pluralism over pluralistic naïve realism and no considerations running the other way.

Appendix 6F. Flexible naïve realism

Here I compare pluralism to *flexible naïve realism*. As the name suggests, the latter theory is a version of naïve realism. It therefore holds that to perceive is to stand in a primitive relation of awareness, which is in this case dubbed *presentation*, to one or more instances of sensory qualities. However, the flexible naïve realist also holds that to *hallucinate* is to be presented with one or more instances of sensory qualities. How is this combination of views possible? It is because sensory qualities are *ontologically flexible*: the existence of an instance of redness can be fully grounded *either* in its inhering in an ordinary object such as a tomato, *or* in a subject's being presented with it.

Thus, the flexible naïve realist says that when I hallucinate a red tomato, I am aware of a perfectly real but mind-dependent instance of redness. The existence of this instance of redness is fully grounded in the fact that I am presented with that instance of redness, and the instance does not inhere in anything – not even my mind or experience. But what happens when I really see a red tomato? Then I am presented with an instance of redness whose existence is *ontologically over-determined*: its existence is fully grounded in the fact that I am presented with it *and* in the fact that it inheres in the tomato.⁴¹

Flexible naïve realism has many virtues. It accommodates the attractive idea that perceptions of redness and hallucinations of redness are deeply similar, for it says that both involve genuine awareness of instances of redness. It honors the powerful intuition that we are *presented* with instances of redness, even when we are hallucinating. At the same time it retains

⁴¹ Flexible naïve realism is carefully developed in a series of papers by Umrao Sethi. See her (2020); (forthcoming a); (forthcoming b).

the natural idea that tomatoes can be red, for it says that redness is ontologically flexible and thus capable of inhering in tomatoes. It even accommodates the idea that the tomato that I see is *mind-independently* red. For that very instance of redness could (and likely will) continue to exist whether or not there is anyone to perceive it. Its inherence in the tomato is enough to guarantee its existence.

I believe that something brilliant in the way that flexible naïve realism is constructed. I prefer it to standard naïve realism. Still, I think the theory faces several problems – a superficial one and some deeper ones.

An initial problem is that there is nothing in flexible naïve realism to explain why seeing a tomato reveals many essential truths about redness. The flexible naïve realist might respond by adopting an element of pluralism: she might say that being presented with something reveals a substantial part of its essence. But that creates a further problem. Presumably the flexible naïve realist wants to say that we are presented with *tomatoes*, not just instances of redness – but seeing a tomato reveals nothing of its essence. Still, perhaps this problem can be met by saying that we are presented with tomatoes only in some secondary sense, in virtue of being presented with their instances of redness and roundness.

But there is a deeper problem. As with most naïve realist views, a central motivation for flexible naïve realism is to accommodate the apparent presentational character of perceptions and hallucinations.⁴² I have already argued that this gives the flexible naïve realist no comparative advantage over pluralism. The pluralist offers a fully satisfying theory of apparent presentational character, too.

⁴² See Sethi (2020, pp. 588-589).

In fact, matters are worse than that. The flexible naïve realist does not fully explain why typical perceptual experiences first-personally seem to present their targets. She *does* say that they really do present their targets, but that is not a complete explanation: we need to be told how subjects manage to learn this about them. The obvious answer is that perceptual experiences wear their status as presentational on their sleeves – but this answer runs headlong into the arguments of chapter 5. There I gave examples of perceptions that do not first-personally seem to present anything. I do not know how the flexible naïve realist can explain why certain experiences seem presentational and why others do not without adopting the categorization account; and if she does adopt that account, it is not clear what work is being done by the distinctive commitments of her theory.

I believe that flexible naïve realism faces another deep problem – this time, one that is methodological. The flexible naïve realist develops a theory of the nature of perceptions, hallucinations, and sensory qualities. The commitments of this theory are extremely elaborate. I am willing to accept elaborate commitments *if there is specific evidence that favors them over rival commitments*. But flexible naïve realism is deliberately constructed so as to remove the possibility of such evidence.

For instance, you might have thought to test flexible naïve realism by considering what would happen in the following case. A subject sees a tomato; the tomato is destroyed; and at that very moment, the subject's brain is artificially stimulated to produce an experience of redness. Flexible naïve realism would seem to predict that a single instance of redness would continue to exist throughout this story. For if the existence of the instance of redness is ontologically overdetermined, then even when *one* of its ontological supports, the tomato, is taken away, the instance can persist due to the presence of its *other* ontological support, the brain state. Moreover,

since the original instance of redness was publicly visible, it would presumably continue to be visible. *That* would certainly be a distinctive prediction.⁴³

In fact, however, the flexible naïve realist avoids making this prediction. She says that the subject enters a new brain state as soon as the tomato is destroyed: the subject's brain state when she sees a tomato is individuated by *that very tomato*. Thus, once the tomato is destroyed, the subject enters a qualitatively identical but numerically identical brain state. The result is that even though the instance of redness that inheres in the tomato is ontologically flexible – even though the existence of *that very instance* is fully grounded in the subject's being presented with it – it is *metaphysically impossible* for that instance to exist in the absence of the tomato, purely in virtue of the subject's being presented with it.⁴⁴ (This is, by the way, yet another elaborate commitment of the view.)

This is a conspiracy theory: when the theory *would* deliver distinctive predictions, further commitments are added to cancel those predictions. It is possible that there are metaphysical conspiracies of this sort. But we would need a lot of evidence before we could rationally believe in them. Pluralism, by contrast, makes a number of bold and specific predictions. The correctness of these predictions gives us strong reason to favor pluralism over other theories – including flexible naïve realism.

⁴³ Incidentally, this reveals another problem for flexible naïve realism. Suppose that I hallucinate a red tomato. The flexible naïve realist says that I am presented with a real, albeit mind-dependent, instance of redness. It lies in the nature of redness that its instances are located in space, so where is *this* instance of redness located? She cannot say that it is located in ordinary space: otherwise it would be publicly visible. Thus she must say that it is located in some extraordinary space, such as a mental space. This is a very costly commitment.

⁴⁴ *Ibid*, pp. 606-609.