

5. Hallucination

In this chapter, I focus on hallucination, although the discussion sometimes spills over to perception, perceptual experience, and sensation. I begin by offering an account of apparent presentational character (§1), which leads me directly to an account of the nature of hallucination and sensation (§2). I round out this discussion by considering some epistemological questions about our beliefs about presentation (§3). After a brief interlude, (§4), I then lay out several arguments from hallucination: a basic argument from hallucination (§5), a screening-off argument (§6), and an argument from first-personal access (§7). I offer pluralist rebuttals of each.

1. The categorization account of apparent presentational character

The previous chapter offered a pluralist account of *real* presentational character. But the pluralist also owes an account of *apparent* presentational character – an account of why particular experiences *first-personally seem* to be presentational or non-presentational (as the case may be). If I see a red tomato in mundane circumstances, then it will first-personally seem to me that I am presented with the tomato. It will also seem to me, incorrectly, that I am presented with “the” tomato when I have a perfectly matching hallucination as of a red tomato. More generally:

Apparent Symmetry: Any typical perceptual experience has *apparent presentational character* – that is, it first-personally seems to present one or more targets. Typically, these targets seem to be ordinary, mind-independent objects or events, along with their property-instances.

At the same time, when I experience phosphenes or vividly imagine a tomato, it will not first-personally seem to me that I am presented with anything. Any theorist of perception owes an account of such facts.

One tempting initial thought is that any perception wears its presentational status on its sleeve: if you are presented with a real, spatiotemporally located entity, then it will first-personally seem to you that this is so (if you are rational, attentive, etc.). Call this the *easy account*. Perhaps the account could be extended to sensations, sensory imaginings, and episodic memories: perhaps these wear their non-presentational status on their sleeves, too. We would still need to explain why hallucinations incorrectly seem to present ordinary, mind-independent particulars. But the easy account is a start.

Unfortunately, the account is incorrect. Suppose that, unbeknownst to me, you set up faint flashing lights within my field of vision. When I see them, my experience is indiscriminable from the phosphene experience that I get when I vigorously rub my eyes. This is already problematic for the easy account: arguably, it will not first-personally seem to me that I am presented with real lights rather than just experiencing phosphenes.

We can modify the case to make the point even sharper. Suppose that I vigorously rub my eyes just before I see the faint flashing lights. I simultaneously experience phosphenes and real lights. Since I do not know anything about the lights beforehand, I will take them to be phosphenes. It will first-personally seem to me that I am not presented with them. And it is easy to multiply examples of this sort: for many sensations, there are possible perceptions that first-personally seem to match them precisely. A vivid sensation of a bright green after-image might perfectly match a perception of a bright green light; a sensation of ringing in the ears might perfectly match an actual ringing sound; a sensation of bitterness on the tongue might perfectly match a perception of a bitter taste. The lesson is that some perceptions do not wear their

presentational status on their sleeves: even though they in fact present their targets, they do not first-personally seem to do so.

What account can the pluralist give to replace the easy account? I suggest:

The *categorization account*: For an experience to first-personally seem to present a target is for the subject to subpersonally categorize the experience as a phenomenally conscious perception. In normal, mature human beings, the categorization of an experience as a perception (whether phenomenally conscious or not) occurs when and because the subject's perceptual system detects that there is a distinctive *perceptual gestalt* associated with the experience of the target.

What is it for an experience to be *categorized as a perception*? Focus first on normal human beings. The idea is that, as we mature, we develop a theory-laden concept of what perceptions are.¹ For example, we implicitly appreciate that perceptions involve a relation of direct awareness to a target, a relation that is perspectival and that updates in real time.² Our subpersonal systems then automatically apply this concept to any experience of ours that has a perceptual gestalt. When that occurs, the targets of the experience are taken to really exist. This is what occurs in a typical perceptual experience and does not occur in a typical experience of a mere sensation. This idea can then be extended to creatures without sophisticated conceptual capacities. In such creatures, categorization consists in the application of proto-concepts. These might be very simple indeed, consisting in nothing more than unsophisticated behavioral dispositions.

¹ Though there are possible creatures that form the same concept in a very different way.

² For further suggestions of what might be part of this concept, as well as evidence that such a concept is formed as we mature, see Andrews (2020, esp. pp. 236-239).

I said that in normal, mature human beings, experiences are categorized as perceptions if they have a *perceptual gestalt*. What is that? The idea is that perceiving is characteristically associated with sensory qualities of a certain distinctive kind. For example, when you see a tomato under normal circumstances, there are many cues that you are *perceiving* the tomato rather than sensorily imagining or episodically remembering it.

Some of these cues primarily involve the sensory qualities *of the tomato*: the red* of the tomato is experienced as especially vivid; it is experienced as relatively stable; and it appears to be part of a continuous three-dimensional surface. Then there are other sensory cues that involve the *interaction* between the sensory qualities of the tomato and the sensory qualities of your environment or body. For instance, you may lose your awareness of the redness* of the tomato as the lights dim. In addition, your awareness of redness* will persist only while you are looking in the right direction with your eyes open, and there will be changes in the shape* that you are deeply aware of as you move around the tomato. One last important factor is that your awareness of these qualities will seem to be *involuntary*.³ Note well, though: this does not suffice for a presentational gestalt. Your experiences of gently floating phosphenes or of bright after-images will also seem to be involuntary, but these experiences will not have a presentational gestalt.

The pluralist thinks that no single cue is necessary or sufficient for the first-personal appearance that you are presented with the tomato. Instead, the appearance of presentation occurs whenever you recognize a *gestalt* of these cues – a gestalt that normally occurs only when you are perceiving the tomato. This recognition is not the result of deliberate inference. It occurs

³ Thus a perceptual gestalt is primarily, but not purely, a matter of how the sensory qualities are: involuntariness is an example of a further factor.

automatically and subpersonally, and it is made possible by the perceptual learning that you have undergone in the past.

I submit that the categorization account explains a wide range of data.

Datum. Some perceptions – experiences that really present their targets – do not first-personally seem to present anything. One example is the case described above in which a subject sees faint lights but takes herself to be experiencing phosphenes. For another example, in one famous experiment, subjects looked at a white wall and were asked to imagine a blue banana rotating slowly on the wall. Unbeknownst to them, the experimenters then projected a dim image of a rotating blue banana on the wall. Subjects did not realize that they were seeing a banana; they continued to believe that they were merely sensorily imagining the banana.⁴

The categorization account explains this, for it allows that perceptions may sometimes lack a perceptual gestalt. In the case in which the subject sees faint flashing lights, her perceptions are neither vivid nor well-integrated with the environment. Similarly, in the experiments just described, the subjects' perceptions are not at all vivid, and the targets of their perceptions are unusual: we do not usually come across faint projections on walls.

Some will worry that a subject of this experiment *can* tell that their experience is a perception, since such an experience is involuntary. I address this concern in Appendix 5A.

Datum. Sometimes the first-personal appearances are silent on whether an experience presents anything: I might experience a ringing quality, ringing*, and wonder whether I am really hearing it or just experiencing a ringing in my ears.

With the categorization account in hand, we can explain this by observing that there are gestalts that are neither clearly presentational nor clearly non-presentational. In the present case,

⁴ See Perky (1910).

perhaps the ringing* is not especially vivid and is not coming from any obvious direction. This explanation is further confirmed by facts about what subsequent events might tip my experience into seeming presentational or non-presentational. If I go to investigate and notice that the ringing* increases steadily in volume* as I walk in one direction, then my experience will first-personally seem presentational. If there are no changes in volume* no matter where I walk, that will likely tip my experience into seeming non-presentational.

Datum. It has been said that for a mental event to have apparent presentational character *just is* for it to be a sensory representation of a determinate content.⁵ However, there are counterexamples to this claim from both directions. Suppose that a boxer takes a hard punch to the head and experiences bright, vivid flashes of light. His experience will not seem to be presenting actual flashes of light – it will not seem to him that he is seeing, say, the flashes of the spectators’ cameras. Or suppose that you close your eyes and then open them as little as possible to see a cup in front of you. You will barely be able to make out the cup, and your experience will be very faint. Still, it will first-personally seem to you that you are presented with the cup.

The categorization account predicts this. For it does not appeal to any single property of sensory qualities. It appeals to a perceptual gestalt, where gestalts are understood holistically. In

⁵ This is the view of Pautz (2007, p. 519, emphases added): “Why then do experiences, including hallucinatory experiences, have a presentational phenomenology while thoughts do not? ... My answer to this question is that there is no answer. Some intentional states (believing contents, desiring contents) do not have a presentational phenomenology. Other intentional states do have a presentational phenomenology: states in which one stands in the special “sensory” relation to contents (propositions or complex properties) which I have dubbed ‘sensorily entertaining’. (Or at least some episodes of standing in this sensory relation to contents have a presentational phenomenology; perhaps in imaginings and dreams we stand in the same relation to contents but *because the contents are quite indeterminate* the experiences lack a presentational phenomenology.) The state of having an experience as of the vivid presence of a red thing at a certain place just is the state of standing in this special relation to a *detailed* content according to which there is a red thing at a certain place.”

the case of the boxer, before he experiences vivid flashes, he feels the sensory quality of a sharp blow to his head beforehand, and there is no particular reason for him to see real lights. That explains why his experience of the lights does not seem to him to be presentational. But when you faintly see the cup, you can feel the sensory qualities of your eyelids that indicate that your eyes are not fully open, so you are not surprised to find that your visual experience is faint – not to mention that your experience of the sensory qualities of all other seen objects is faint, too. On top of that, your experience of the cup has many other sensory similarities to a standard perception of a cup: for instance, you experience the sensory qualities indicative of a table supporting the cup. That is why this experience will strike you as presentational.

Datum. It is possible for the first-personal appearances regarding presentation to conflict with what you know. You might know that you are about to see faint colored lights that look just like afterimages and phosphenes. Still, if the timing of the lights is precise enough – if they go off just when you rub your eyes and fade shortly afterwards – then it will seem to you that you are not presented with real lights. Or you might know that your brain is about to be stimulated to create a perfect hallucination of seeing a tree. The hallucination, when it occurs, will still first-personally strike you as presentational.

The categorization account explains this by saying that it is the subject's *perceptual system* that detects whether or not there is a perceptual gestalt. Perceptual systems are famously insensitive to at least some of the subject's knowledge: the Müller-Lyer illusion does not end simply because the subject knows that it is an illusion. This is compatible with perceptual learning, which occurs over extended periods of time.

This makes it clear how the pluralist should handle an important phenomenological objection. The objection is that according to the categorization account, the first-personal appearance of presentation is phenomenologically indirect: you identify your deep awareness as

having a certain gestalt, and from that you arrive, whether by inference or by some other transition, at the first-personal appearance of presentation. But this is plainly incorrect: the first-personal appearance is phenomenologically direct. It just immediately seems to you that you are presented with the tomato.

I agree with my opponent that the first-personal appearance is phenomenologically direct. That is precisely what the categorization account says: it is part of the view that your identification of the perceptual gestalt occurs *subpersonally*, in a way that is beyond your first-personal access. That is why the first-personal appearance of presentation just shows up, with no particular effort on the subject's part.

Datum. Return now to:

Apparent Symmetry: Any typical perceptual experience has apparent presentational character – that is, it first-personally seems to present one or more targets. Typically, these targets seem to be ordinary, mind-independent objects or events, along with their property-instances.

The pluralist easily explains this claim. For the categorization account says that all first-personal appearances of presentation in normal, mature human beings are explained by a particular gestalt. This gestalt is found in any typical perception. To this, the pluralist adds that for any perception, there is a matching hallucination, which she analyzes as a hallucination that involves deep awareness of exactly the same sensory qualities. Thus, the perception and the matching hallucination will have exactly the same gestalts. As a result, the hallucination will also first-personally seem to present various entities – in particular, ordinary, mind-independent particulars of the same sort that seem to be presented in the perception.

In the appendices, I discuss some further data that support the categorization account. The first datum is that certain fantastical experiences, such as prophetic experiences of the future, first-personally seem to be presentational (Appendix 5B). The second datum is that there are pairs of possible perceptions that are sensorily identical but that differ in whether they first-personally seem to present (Appendix 5C).

Conclusion. The data concerning apparent presentational character turn out to be much more complicated than you might expect. But they are neatly explained by the categorization account.

The data support a further conclusion. Distinguish between two kinds of representations. For a representation to be *committal* is for the subject, simply in virtue of having that representation, to at some level take things to be as they are represented. For a representation to be *non-committal* is for this not to be the case. Beliefs are paradigms of committal representations: in virtue of believing that p , a subject must take p to be the case *at some level* – even if, at the same time, she holds the inconsistent belief that $\text{not-}p$. Desires are paradigms of non-committal representations: wanting to eat ice cream now does not make it the case that you take yourself to be eating ice cream now.

The standard representationalist view is that all perceptual experiences – unlike, say, sensory imaginings – are committal representations.⁶ We can now appreciate why this is false:

⁶ For instance, Pautz says, “On Intentionalism, experiences are a type of intentional state akin to thoughts” (2007, p. 519); Siegel says, “perception involves states that are importantly similar to beliefs: visual experiences” (2010, p. 27); Millar speaks of “the sort of commitment that is present in belief and perceptual experience and absent from merely grasping or considering some proposition” (2014, p. 250). The same view is implicit but clear in Tye (1995, p. 142). The point has been given fuller expression by naïve realist Michael Martin: “the only plausible forms of intentional theory [i.e., representationalism] appeal to the stative [i.e., committal] notion of representation in order to explain the distinctive phenomenology of perceptual experiences”

there are unusual perceptions, such as perceptions of faint blue bananas or of dim lights after the subject has rubbed her eyes, that are not categorized, via their gestalts, as perceptions. The subject does not take the contents of these perceptions to be true at any level. That is why the pluralist instead says that *all* perceptual experiences are non-committal representations – even when the subject *does* categorize the experience as a phenomenally conscious perception. For even then, the subject does not take the content of her experience to obtain *merely because she has the experience*. She takes the content of her experience to obtain *because she categorizes it in a certain way*.

2. Hallucination and sensation

The pluralist says that to be a perception is to be a mental event consisting of a perceptual relation to a target. The perception is access-conscious if the perceptual relation is “lit up” by deep awareness (i.e., it appropriately causes deep awareness to be applied to the target) and access-unconscious otherwise. But what might the pluralist say about hallucination and sensation?

The pluralist’s theory is inspired by the categorization account. She says this:

To be a *hallucination* is to be an involuntary sensory representation that is not a perception but is subpersonally categorized as one. In normal, mature human beings, this occurs when and because the content of the sensory representation is taken to have a perceptual gestalt.

(2002, p. 387). Martin makes it clear that the phenomenology that he has in mind is apparent presentational character (pp. 387-388).

To be a *sensation* is to be an involuntary sensory representation that is not a perception and is subpersonally categorized as *not* being a perception. In normal, mature human beings, this occurs when and because the content of the sensory representation is taken not to have a perceptual gestalt.

According to this theory, 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. Finally, what distinguishes hallucinations from sensations is whether the content of the sensory representation is taken to have a perceptual gestalt.

On that last point, note that what distinguishes hallucinations from sensations is not whether the experience *actually has* a perceptual gestalt. What matters is whether it is *taken* to have one by the subject's perceptual system. 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 small floating lights. Such an experience is clearly a hallucination, not a sensation. (In Appendix 5D I briefly consider what to say about instances of involuntary deep awareness that the subject's perceptual system does not categorize either as having or lacking a perceptual gestalt.)

But there is an obvious objection to the theory. 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 it as a hallucination or sensation. But the experience is actually a sensory imagining. 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.

The pluralist can reply 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.*⁷ 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.)

The pluralist's theory of hallucination and sensation has many virtues. We have already seen one of them: the theory lets us explain why matching hallucinations will, in the usual case, first-personally seem to present ordinary, mind-independent particulars. Again, the explanation is that a perception and a matching hallucination will involve deep awareness of precisely the same profile of sensory qualities. Thus both will be subpersonally categorized as having a perceptual gestalt, and in either case the subject will take herself to be presented with targets 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.

Likewise, this theory makes it easy to explain why a phenomenally conscious perception and a matching hallucination would cause any subject to form similar beliefs and take similar

⁷ For an alternative response to an objection in this ballpark, see Ichikawa (2009, p. 107).

actions.⁸ This occurs because each experience involves deep awareness of precisely the same sensory qualities and is consequently subpersonally categorized in the same way.

Another virtue of the theory is that it allows for the possibility of phenomenally unconscious hallucinations and sensations. After all, in a normal mature human being, an involuntary sensory representation will be subpersonally categorized as having a perceptual 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 sensations.⁹ There is an apparent obstacle to such an explanation: the pluralist analyzes hallucinations and sensations partly in terms of categorization, which might seem to occur only in a creature with full-fledged conceptual capacities. However, recall that the pluralist is working with a more generous notion of categorization. She says that in simple creatures, categorization might consist merely in the application of proto-concepts, which in the simplest case might amount to nothing more than the formation of certain behavioral dispositions. In this way the obstacle is removed.

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.

⁸ Assuming, of course, that we hold fixed her background mental situation.

⁹ It is notoriously difficult for naïve realists to explain this datum.

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.¹⁰ More importantly, there might be a psychological kind formed by all and only the experiences that first-personally seem to be presentational. 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 seen that some perceptions, even though they are in fact presentational, do not first-personally seem to be presentational. 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: all perceptual experiences are instances of deep awareness. However, this kind also includes non-perceptual experiences such as sensations, episodic memories, and sensory imaginings. In this sense the pluralist accepts *some* sort of common kind view, albeit one that is weaker than even the weak common kind view.

¹⁰ Not even properties, since the relation of presentation must be perspectival and must update in real time.

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 something. 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 something: again, there are perceptions that do not first-personally seem to present. In this way, the pluralist undercuts the very motivations for the weak common kind view.

Now for other matters. Consider some views about perception:

- (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.

These views are motivated by the desire to distinguish perceptions from sensations. In Appendix 5E, however, I argue that these views draw the distinction in the wrong place. Pluralism does better.

Finally, it is worthwhile to consider other kinds of experiences, such as sensory imaginings and episodic memories. What is the nature of these experiences? In addition, such

experiences usually first-personally seem to be the kinds of experiences that they are – but not always. What explains these first-personal appearances? I conjecture that these questions can be answered by extending the pluralist ideas developed in this chapter. See Appendix 5F.

I move now to some epistemological matters.

3. The epistemology of presentation

We sometimes form beliefs that we might express by saying, “I am *presented* with that thing.” Call any such belief a *presentational belief*. Can the pluralist provide a satisfying epistemological account of such beliefs? Suppose for example that I see a blue ball in front of me. I come to believe, in the usual way, that I am presented with that thing. My belief amounts to knowledge. Now suppose that I suffer a matching hallucination and form a similar belief: the belief that I am presented with “that thing.” My belief, though it obviously does not amount to knowledge, is evidently justified. How can the pluralist explain these facts?

It is easy enough to see why my presentational belief would amount to knowledge when I am genuinely perceiving: my belief meets all of the standard conditions on knowledge. It is true, of course. It is also *safe*. Consider all nearby worlds in which I am presented with that same blue ball and form a presentational belief on the same basis. In such worlds, I might see the ball from a slightly different vantage point. My subpersonal systems will recognize that my experience has a perceptual gestalt and categorize it as a phenomenally conscious perception, and on that basis I will form the belief that I am presented with that thing. My belief will be true in all such nearby worlds.

My presentational belief will even be *sensitive*: in the nearest worlds in which I am *not* presented with that thing – say, worlds in which I turn my back on the ball or simply close my

eyes – I will not believe that I am presented with that thing. In addition, my presentational belief will be *formed via a reliable process*. For consider any process for forming beliefs on contingent matters. If the process is both safe and sensitive, it is hard to see how it could be unreliable.¹¹

Perhaps some will object that when I have a matching hallucination, I will form a similar presentational belief.¹² But, on pain of skepticism, this cannot undermine the status of my presentational belief as knowledge in the case of genuine perception. Perception surely positions me to know that I have hands, even though there are possible scenarios in which I form a similar but non-knowledgeable belief on the basis of hallucination. There is no evident reason that my belief that I am presented with that thing would be epistemically worse off than my belief that I have hands.

Now turn to the presentational belief that I form when I merely hallucinate a blue ball. This belief does not amount to knowledge, of course. It is not even true. But it is plausible that the belief is justified. Can the pluralist explain why? One natural answer is that the hallucinatory presentational belief is the output of a reliable belief-forming process. Another natural answer is that a belief is justified as long as it is simply the endorsement of how things first-personally seem to the subject, in the absence of defeaters. The hallucinatory presentational belief would meet this condition, as well.

I conclude that the pluralist has a simple and unproblematic account of the epistemology of presentation.

¹¹ It is important to include the restriction to contingent beliefs. Necessarily true beliefs might trivially turn out to be safe and sensitive.

¹² Not quite the *same* belief, since my hallucinatory presentational belief will have only failed singular content.

4. 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 manipulation 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 distinguish and rebut three arguments against pluralism that focus on these subjects: a basic argument from hallucination (§5), a screening-off argument (§6), and an argument from first-personal access (§7). Along the way, I reveal some important advantages of pluralism over standard naïve realism.

5. 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.¹³

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.

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.¹⁴ 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. Of course, the pluralist provides such an explanation: it is that both experiences involve deep awareness of precisely the same sensory qualities. Given the

¹³ Cite Robinson, Foster, etc.

¹⁴ On this point, see Langsam (1995, pp. 43-44); Martin (2004, pp. 55-58).

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 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.¹⁵ 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.¹⁶

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

¹⁵ Again, see fn. 2.

¹⁶ 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).

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.

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 both theoretically simpler and truer to the first-personal appearances.

6. A screening-off argument

Naïve realism is famously subject to the following argument¹⁷:

The *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.

¹⁷ 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 (SO3) by accepting certain claims about fundamental kinds (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 weak 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.

- 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.

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 argument¹⁸:

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

¹⁸ 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.

- 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.

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 universals. She even thinks that the 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. For 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 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.¹⁹

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 sensory qualities. As a consequence, (ii) both experiences have a perceptual gestalt. Because of this, (iii) they will both first-personally seem to present 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.²⁰

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 red*. Presentation is a form of consciousness! It brings the tomato and its instances of red* into the subject's conscious life. For the pluralist, the perceptual kind is crucial to the explanation of

¹⁹ 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.

²⁰ Again, however, the hallucinator's judgments will be only putatively singular.

this aspect of the perceiver's conscious perspective. There is no such aspect to the conscious perspective of the hallucinator.²¹

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

7. 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.²²

²¹ Here I am understanding consciousness as *awareness*. You might instead understand consciousness as *phenomenal consciousness*, which is in turn understood in terms of *hard character*. In that case the pluralist would say that the tomato is not an aspect of the subject's consciousness. As I argued in ch. 2, this is independently plausible: there are no hard problems of consciousness associated with being presented with tomatoes.

²² This line of argument is pressed in Martin (2004, pp. 59-60).

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.²³ 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 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.²⁴ 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.²⁵

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.²⁶ 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 red* and round*.

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

²³ My opponent might want to modify the argument to refer to mental kinds that are also *fundamental*. This would not affect my response.

²⁴ See Martin (2006, p. 369).

²⁵ 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.

²⁶ Cite Martin and others.

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. 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.²⁷

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.²⁸

²⁷ 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. Thus I prefer the pluralist account.

²⁸ See Martin (2004, pp. 53-58) and (2006, p. 369).

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.

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.²⁹

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 red*. 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.³⁰ 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

²⁹ See Martin (2004, pp. 59-60).

³⁰ 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.

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 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.

8. Conclusion

It is worth seeing how pluralism can be inserted into the debate between representationalism and naïve realism.

One of the most significant challenges to representationalism is to explain why perceptual experiences have apparent presentational character. What is difficult is to identify precisely how perceptual experiences differ from thoughts, given that these are both representations, and from sensory imaginings and episodic memories, given that these are both sensory representations. The pluralist meets this challenge: she says that apparent presentational character is a matter of deep awareness that is categorized as having a perceptual gestalt. Sensory imaginings and episodic memories do not typically involve deep awareness with such a gestalt, and thought does

not constitutively involve deep awareness at all. Thus there is no apparent presentational character in these cases.

At the same time, pluralism is an improvement on standard naïve realism. Pluralism has the resources to block the basic argument from hallucination, the screening-off argument, and the argument from first-personal appearances. Perhaps standard 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.

Thus pluralism is a deeply satisfying theory: it fully honors all of the key intuitions about perception, hallucination, and the relationship between the two.

It is natural to worry that the pluralist does not *fully* honor the intuition behind standard representationalism. The intuition is not just any perception and any matching hallucination have *a* deep commonality; it is that they are *fundamentally the same*. Or so the worry goes.

I agree that this 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 cups, trees, and tables to us. They also position us to make singular reference to novel particulars, and to know what is thereby being referred 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.

The result is an account that is beautifully unified in many respects. There seems to be something that unites all sensory mental events, whether phenomenally conscious or not: the

pluralist explains this in terms of sensory representation. There seems to be something that unites all phenomenally 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 more specific that unites all mental events with apparent presentational character, including typical perceptions and hallucinations but excluding sensations: the pluralist explains this in terms of categorization via a perceptual gestalt. Yet there also seems to be something distinctive to phenomenally conscious perceptions: the pluralist says that this is a matter of a particular form of representation that is lit up with deep awareness.

I conclude that pluralism is the best theory of perception to date. The next chapter mops up by explaining why perception positions us to conceive of objects and properties as categorical, intrinsic, mind-independent, and so on.

Appendix 5A. A concern about the Perky experiments

As part of my argument against the easy account, I appealed to experiments in which subjects mistook *perceptions* of faint images of blue bananas for *sensory imaginings*. I suggested that some experiences like this do not first-personally seem to present, though in fact they do present.

You might challenge the datum. You might say that some subjects did not notice that their experiences were presentational, but had they been suitably attentive, they could have. In particular, they could have noticed that their experiences were involuntary. After all, when a subject sensorily imagines a blue banana, she can influence the shape, position, and color of the

experienced banana simply by willing it. A subject who sees an image of a blue banana cannot do that.³¹

Suppose that this is correct. Still, the case can be modified to avoid this problem.

As before, consider a subject who is instructed to look at a white wall and imagine a blue banana on it. This time, however, the subject is told that scientists will sometimes be projecting a blue banana image on the wall. Her task is to determine whether each of her blue banana experiences is a perception or a sensory imagining. Suppose that the experiment is run by future neuroscientists who can detect, through advanced neuroimaging, exactly when the subject forms the intention to imagine a blue banana. Our scientists can also predict the exact shape, orientation, color, and so on that the imagined banana would have. Whenever the subject is about to imagine a blue banana, the scientists they project a faint image of a blue banana on the wall that has exactly the right shape, orientation, and color. The projected image will even rotate, change color, and so on in accordance with the subject's intentions.

A subject of this experiment would mistakenly take her perceptions to be voluntary. These perceptions would not first-personally seem to present anything, even though they would in fact present blue-banana images. This restores the counterexample to the easy account.

Appendix 5B. Fantastical experiences

As another line of support for the categorization account, I wish to mention a pair of fantastical experiences that first-personally seem to present their targets.

³¹ I thank Matthew Walker for pressing this objection.

Suppose that Harry Potter is using the pensieve, a magical device which lets him experience any memories that have been stored there. When the pensieve is activated, it is as though Harry is actually reliving the remembered event: he seems to have a ghostly body that he can use to explore the remembered area, and it is as though he is really seeing, hearing, and smelling the relevant past events. Or imagine that Frodo Baggins is gazing at the Mirror of Galadriel, a magical artifact that lets him experience events across space and time. The mirror shows him events from the past, present, and future, but it does not indicate which of these time periods the event is from. When an event is revealed to Frodo, his experience has the same vivid, stable character as his ordinary perceptions.

It is obvious that Harry's pensieve experiences and Frodo's mirror experiences would first-personally seem to be presentational. The categorization account predicts this, given that these experiences have a perceptual gestalt. That is one more mark in favor of the account.

What about *real* presentational character – do these fantastical experiences have that? I myself have no strong pre-theoretic intuitions about the matter, so I am content to defer to our best theory. Pluralism entails that these experiences do not have real presentational character, since neither experience *really* conveys information about the spatial location of the target relative to the subject. I find that verdict plausible enough.

Appendix 5C. Another datum that favors the categorization account

Let me take some time to build up to this last datum.

Imagine that you seem to see a tomato in the usual way. You reach to pick it up – but your hand passes through the tomato, as though it were a hologram! Astonished, you try several times more, with no better luck. You try to prod the tomato with a fork, but that, too, passes

right through. When you mention this astonishing phenomenon to others, you discover that they do not seem to see the tomato at all. You head to the market to investigate more tomatoes and seem to find a bin full of them. Again, your visual experience seems to be mundane. But when you try to touch the tomatoes, your hands pass through them once more, and the clerk whom you flag down for help tells you that the bin is empty.

If these sorts of experiences persist for long enough, then mundane tomato-experiences might cease to first-personally seem to present you with tomatoes. You might take yourself to be having mere sensations – tomato-phosphenes, perhaps.³²

Suppose that all of this was an elaborate prank: you were seeing perfectly ordinary tomatoes all along. A mischievous demon made parts of the tomato vanish as you reached for them, while counteracting the effects of gravity on any unsupported tomato parts that remained. She then replaced the missing tomato parts whenever you removed your hands. She also tampered with the visual experiences of those around you, temporarily blinding them to the existence of the tomatoes around them.

Now to the point: compare two of your experiences. First, your experience of a tomato before the antics of the demon. Second, your experience of a tomato afterwards, once you take all such experiences to be of tomato-phosphenes. Both of these experiences are perceptions of real tomatoes. They might even be sensorily identical: you might be looking at visually indiscriminable tomatoes from the same perspective in the same lighting. Still, only your first tomato-experience will first-personally seem to be presentational. Your second experience will not. The datum, then, is that there are pairs of possible perceptions that are sensorily identical but that differ in whether they first-personally seem to present.

³² So far, this case is structurally identical to the doll case described in Siegel (2006, pp. 369-370).

This is somewhat surprising! The categorization account explains why it is possible. The point is that we subpersonally categorize experiences as perceptions on the basis of (apparent) perceptual learning; it is perceptual learning that leads us to conclusions about which sensory gestalts are presentational and which are not. Thus experiences involving precisely the same sensory qualities can be categorized as presentational or not, depending on what (apparent) perceptual learning the subject has undergone.

Appendix 5D. Arguments against some views about perception

The data that support the categorization account 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>.³³ This content can be gussied up further if desired, for example by saying that the causation is automatic, linked to the present state of the object, and not mediated by any distinct experience.³⁴ 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.

³³ This theory is introduced in Searle (1983, ch. 2).

³⁴ This is how the theory is developed in Millar (2014a, esp. p. 249) and (2014b, esp. pp. 642-643).

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. But 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 experience: I would experience whatever is on the other side of the vase. Call this the *perspectival content view*.³⁵

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.³⁶ The thought is that when I have a sensation of a bright green after-image, I do not take the after-

³⁵ This view is defended at length in Siegel (2006).

³⁶ Ibid, p. 355.

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 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.³⁷

³⁷ See Siegel (2006, pp. 369-370).

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 *content* of her experience. The categorization account instead analyzes this in terms of the subject's *subpersonal categorization* of her experience.

The categorization account is thereby able to enjoy the advantages of these alternative views. In the main body of this chapter, I explained in detail how it can capture the difference between perceptual experiences and sensations. Likewise, the categorization account 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 your subpersonal systems cease to categorize your

doll-experiences as having a perceptual 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 perceptual gestalt with respect to other environmental objects.

The categorization 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 experience³⁸ – 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.

There are various replies available to the theorist who posits sophisticated contents of experience.³⁹ But the pluralist avoids the problem altogether. She puts the sophistication into the process by which the subject's *subpersonal systems* categorize her experiences. 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. But there is a third view worth considering: the *dual component theory*, which says that to

³⁸ As Siegel (2010, pp. 42-44) rightly emphasizes, only what is conveyed to the subject is fit to be a content.

³⁹ See Siegel (2006, pp. 382-383); Millar (2014b, pp. 649-650).

be a perception is to be a perceptual belief (or a perceptual seeming) that is caused by a sensation.
(To be added.)

Appendix 5E. Ambiguous experiences

What does the pluralist say about instances of involuntary deep awareness that the subject's perceptual system does not take either as having or as not having a perceptual 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 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, my perceptual system does not categorize 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*-perceptual 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 perceptual gestalt. But, for what it is worth, I prefer the account as written.

Appendix 5F. 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 is appropriately caused by stored perceptual information.

To be a *sensory imagining* is to be a voluntary sensory representation that is not appropriately caused by stored perceptual information.

Note that I am using the term “information” factively. Otherwise the conjectures would be subject to immediate 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. What explains this? It is natural for the pluralist to respond by extending the categorization account. 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 subpersonally categorize the experience in that way. In normal, mature human beings, this occurs when and because the subject’s subpersonal system

detects that there is a distinctive *memorial* or *imaginative profile* associated with the experience of the target.

The idea is that the subject's subpersonal systems attempt to keep track of various facts about a sensory representation. For example: what beliefs, desires, emotions, 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; whether the content of the representation was generated by stored perceptual information; and so on. On this basis, the subject's subpersonal systems arrive at a categorization of each sensory representation as a perception, sensation, episodic memory, or sensory imagining. All of these concepts are theoretical ones, developed as the subject matures.

I have already discussed what cues would indicate that the 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, desires, emotions, 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 stored perceptual information (or apparent information). Alternatively, they might track whether 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 a categorization of the representation as an episodic memory. It will be part of the subject's theoretical understanding of what episodic memory is that its content concerns the subject's 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 subject's subpersonal systems just make an error in processing the cues.

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 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 it, 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. 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. I trust that it is clear how these explanations would go.