

6. Conceptions of properties and objects

Suppose that I see a red, round tomato. My perception positions me to form a certain conception of the *properties* that I perceive. It also positions me to form a certain conception of the *object* that I perceive, the tomato itself. In this chapter, I offer a pluralist explanation of how all of this is possible.

1. Perception, categoricity, and intrinsicity

Begin by distinguishing two kinds of properties. There are *dispositional properties*: salt is disposed to dissolve when placed in warm water, glass is disposed to break when struck. But the instantiation of a dispositional property is ordinarily grounded in the instantiation of one or more properties that are not dispositional, but rather *categorical*. The instantiation of the dispositions just mentioned, for instance, are grounded in the instantiation of certain categorical microstructural properties.¹ The distinction is an obvious one, although, as with many obvious distinctions, it is a difficult question how to characterize it precisely.

There is another obvious but hard-to-characterize distinction: the distinction between *intrinsic* and *extrinsic* properties. Intrinsic properties are properties that concern how an entity and its parts are in and of themselves, regardless of relations to other entities. Extrinsic properties are properties that are not intrinsic. Extrinsic properties should not be confused with *relational properties*, as some relational properties of an entity are intrinsic to it. Take for example

¹ I am thinking of *grounding* as a metaphysical relation that backs metaphysical explanations in much the way that causation backs physical explanations (Schaffer (2016); Wilson (2017)). Its relata are facts, and for p to (partly/fully) ground q is for q to obtain (partly/wholly) in virtue of p .

the property of having more hair on one's head than one's feet. This is an intrinsic property, since it concerns only how the entity in question (a person, say) is. But it is also a relational property: it concerns a relation of quantity between head hair and foot hair.

Categoricity has sometimes been confused with intrinsicity.² This should be avoided. Consider the property, which I currently instantiate, of being less than a meter from my laptop screen. This property is categorical: being less than a meter from my laptop screen is not a matter of having dispositions.³ But it is evidently an extrinsic property of mine, not an intrinsic one.⁴ In addition, it seems that there could also be intrinsic dispositions: dispositions that something has purely in virtue of how it and its parts are. These might include dispositions for the entity's parts to interact with each other in particular ways.

Perception plays an important role with respect to both categorical properties and intrinsic properties. Suppose that you are looking at a round tomato on your kitchen counter. Roundness is a categorical property. It grounds various dispositional properties, such as the property of being disposed to roll down slanted surfaces, but roundness is not to be identified with such properties. Roundness is also an intrinsic property: it concerns how the tomato is, not how it is related to other things. What is interesting is that when you see the tomato, your experience does more than just make you aware of roundness, a property that is *in fact* categorical

² Campbell (2002) consistently elides the distinction. For example, he holds that “[e]xperience is experience of the categorical,” but he calls this the *Intrinsicness Condition* (p. 137). Similarly, he says, “We would ordinarily regard shape properties as the paradigmatic *categorical* properties of objects. Roundness is the reason why the thing tends to roll if suitably propelled and so on; it is the *intrinsic* ground of the complex functional state described” (p. 139, emphasis added). For more discussion of why Campbell might make this mistake, see Appendix 6D.

³ Quite the opposite: instantiating this property might help to *ground* the instantiation of various dispositional properties. Partly in virtue of being less than a meter from my laptop screen, I am disposed to see the screen clearly.

⁴ There is also a semantic distinction between the *concept* of categoricity and the *concept* of intrinsicity: the former is understood in contrast to dispositional properties, the latter in contrast to properties that involve relations to wholly distinct entities.

and intrinsic. Your experience also positions you to appreciate the categoricity and intrinsicity of this property in a special way. More generally, there are two pieces of data here:

The *perceptual categoricity datum*: Perceptions often position us to appreciate, in a special way, the categoricity of certain perceived properties.

The *perceptual intrinsicity datum*: Perceptions often position us to appreciate, in a special way, the intrinsicity of certain perceived properties.⁵

We can get a better grip on these data by considering a foil. Imagine that you do not know what a tomato is, nor do you know what round things are. Someone tells you that there are objects called *tomatoes* that have various dispositions – for example, they are disposed to move down slanted surfaces in a particular way. You are also told that tomatoes have this disposition in virtue of having another non-dispositional property known as *roundness* – a property that concerns how the tomato is in and of itself, regardless of how it is related to other things. You now know *that* the property of roundness is both categorical and intrinsic. But something importantly different happens when you actually see the round tomato: you come to appreciate the categoricity and intrinsicity of roundness in a new, epistemically deep sense.

2. Some standard naïve realist explanations

It has been suggested that the standard naïve realist can give an informative explanation of these data. I disagree. In this section, I will argue that although the standard naïve realist *can* explain

⁵ Campbell identifies both data points, though he runs them together. See fn. 2 and Appendix 6D.

the categoricity and intrinsicality data, her explanation is brute and lacks independent motivation.

The starting point for an informative standard naïve realist explanation will be the claim that the perceiver stands in a primitive, non-representational relation of awareness to the tomato and its instance of roundness, an instance that is in fact categorical and intrinsic. In virtue of this, the tomato and its instance of roundness are *constituents* of the perceiver's experience. The relation of constitution may be understood in many different ways, but this is what is supposed to explain the categoricity and intrinsicality data.⁶

I observe that this explanation has a missing link. Suppose for the sake of argument that the subject's experience is constituted by an instance of roundness that is *in fact* categorical and intrinsic. How exactly is that supposed to position the subject to *appreciate, in a special way*, the categoricity and intrinsicality of the instance of roundness? It is very hard to give a satisfactory answer to this question. For instance, the standard naïve realist might appeal to any of the following principles (where *CI* stands for *categoricity* and *intrinsicality*):

- CI1. The subject is in a position to appreciate, in a special way, any necessary property of the constituents of her experience.
- CI2. The subject is in a position to appreciate, in a special way, any essential property of the constituents of her experience.

⁶ This is what Campbell thinks: "Suppose, however, that we drop the notion that the intrinsic characteristics of experience are *caused* by the categorical objects and properties being seen. Suppose instead we think of the intrinsic characteristics of experience as *constituted* by the categorical objects and properties being seen.... Then we are indeed in a position to understand how the phenomenal content of experience can explain our grasp of what the world is intrinsically like" (2002, p. 156).

CI3. The subject is in a position to appreciate, in a special way, any intrinsic property of the constituents of her experience.

All of these approaches overgenerate. Part of the problem is that the standard naïve realist thinks that *the tomato itself* is also a constituent of the experience. So, if she accepts CI1, CI2, or CI3, she must predict that the subject is in a position to appreciate any necessary property, essential property, or intrinsic property, respectively, of the tomato that she sees. All of these claims are plainly false. Just seeing a tomato does not position you to appreciate its necessary and essential property of being a fruit,⁷ nor does it position you to appreciate the tomato's intrinsic property of having such-and-such mass. Botany and physics are not that easy.⁸

CI1-CI3 overgenerate even if we set aside perceived objects and focus just on perceived property-instances. Some parts of the origin of the instance of roundness – perhaps the time when it began to exist, or its immediate causal antecedents – are arguably necessary and essential to it. But seeing an instance of roundness does not position you to appreciate any of this in a special way, contrary to CI1 and CI2. In addition, as we saw in chapter 2, it is intrinsic to any instance of green* that it is elemental rather than compound. But just seeing an instance of green* does not position you to appreciate its elementality.

In light of these problems, I see only one option for the standard naïve realist: she must say that it is just a *brute fact* about that perceptions often position their subjects to appreciate, in a special way, the categoricity and intrinsicity of certain perceived properties. This will of course explain the data – but the explanation is not informative. Nor is there any independent

⁷ CI1 and CI2 also overgenerate even

⁸ The defender of CI3 might respond that, properly speaking, the subject does not see the entire tomato; she sees only its facing surface, a portion of its skin. This does not help: seeing the facing surface of the tomato does not position you to know the mass of the portion of skin that you see.

reason to accept it: the only purpose of the posit is to explain these data points. A naïve realist who takes this approach will therefore be at a significant disadvantage against a theorist who can offer a properly informative and independently motivated explanation.⁹

The pluralist is in a position to do this.

3. A pluralist explanation

We want to explain:

The *perceptual categoricity datum*: Perceptions often position us to appreciate, in a special way, the categoricity of certain perceived properties.

The *perceptual intrinsicality datum*: Perceptions often position us to appreciate, in a special way, the intrinsicality of certain perceived properties.

According to the pluralist, when a subject sees a round tomato, she deploys a sensory representation that stands in a perceptual relation to the tomato and its instance of roundness. As a result, she has a singular representation of the tomato and its instance of roundness. This representation is lit up with deep awareness: the subject is deeply aware of the quality of roundness, and she takes the instance to be an instance of the quality. In general, deep awareness of any quality positions the subject to know part of the essence of that quality. This makes it easy

⁹ The standard naïve realist will not even have the upper hand against a standard representationalist. For the representationalist can make a similar brute posit, without independent motivation, about how perceptual representation works: she can say that it is just a brute fact that perceptual representations often position their subjects to appreciate, in a special way, the categoricity and intrinsicality of certain perceived properties.

for the pluralist to explain the data, for it is essential to the quality of roundness to be categorical and intrinsic. Thus deep awareness of roundness might position the subject to know these facts.

I emphasize that on this view, deep awareness is a form of awareness-*of*, not awareness-*that*. In particular, it is a form of awareness of a target that positions the subject to know part of the essence of that target. This is why perceptions position us to have appreciation *of a special kind* of the categoricity and intrinsicity of roundness. This is very different from what happens when a subject is simply told that there is a property, roundness, that is both categorical and intrinsic. The latter subject is in a position to know that roundness is categorical and intrinsic, but this knowledge is not grounded in a non-propositional awareness of roundness.

This pluralist explanation has a pair of immediate advantages over the standard naïve realist explanation. We saw earlier that the standard naïve realist must give a brute explanation of the categoricity and intrinsicity data. In particular, she must say that it is just a brute fact about the perceptual relation that it makes the subject aware of the categorical and intrinsic characters of certain properties. In addition, she has no independent motivation for this posit: the only reason for the standard naïve realist to say that the perceptual relation works in this way is to explain the categoricity and intrinsicity data.

Compare this to the pluralist's explanation, which relies on the posit of a relation of deep awareness that reveals part of the essence of its targets. This explanation is not brute. The pluralist takes appreciation of categoricity and intrinsicity to be instances of a more general phenomenon: appreciation of partial essences. Relatedly, this posit was not introduced just to explain the categoricity and intrinsicity data. The pluralist has many independent reasons for positing such a relation: it explains patterns of facts about the hard problems of consciousness, knowledge of essential truths, asymmetries between perceptions of particulars and perceptions of sensory qualities, etc.

So far, then, the pluralist's explanation is very promising. And it has a further explanatory virtue. To appreciate what this is, recall that the pluralist claims that deep awareness occurs in the same way in any perception *and any hallucination that matches it*. The pluralist is now saying that when a subject perceives a round tomato, her deep awareness of roundness positions her to know that this property is categorical and intrinsic. This yields the predictions that a hallucination, just as well as a perception, can position the subject to know that the property of roundness is categorical and intrinsic. More generally, the pluralist makes these predictions:

The *hallucinatory categoricity datum*: Take any perception that positions the subject to appreciate, in a special way, the categoricity of certain perceived properties. A matching hallucination will put the subject in the same position with respect to those same properties.

The *hallucinatory intrinsicity datum*: Take any perception that positions the subject to appreciate, in a special way, the intrinsicity of certain perceived properties. A matching hallucination will put the subject in the same position with respect to those same properties.

I believe that these predictions are correct. For suppose that Beatrice is an unusual subject: she has seen all sorts of objects, but for whatever reason, she has never seen anything round. One day, Beatrice merely *hallucinates* something round – a red tomato, perhaps. (She has encountered plenty of tomatoes before, but they have been roughly cubical due to genetic mutations.) Beatrice's hallucination will not position her to know anything about "the" tomato – there is no such thing. However, the hallucination *will* position Beatrice to appreciate certain

truths about the property of roundness in a special way. In particular, she is in a position to appreciate, in the special way that we have been discussing, the categoricity and intrinsicity of roundness. In this respect, hallucination is on all fours with perception. Had Beatrice been perceiving a round tomato rather than hallucinating one, this would not in any way have improved her ability to appreciate the categoricity and intrinsicity of roundness.

Conclusion. Not only perceptions, but also certain hallucinations, position us to appreciate the categoricity and intrinsicity of certain properties in a special way. The pluralist's theories of perception and hallucination let her explain such facts seamlessly. It has been argued that it is impossible to perceptually represent objects as having categorical and intrinsic properties. The core concern is that perceptual representations are sensitive only to dispositional and extrinsic properties. I respond to several variations of this concern in Appendix 6A.

Now consider some competing accounts of hallucination that have been advanced by naïve realists:

- (i) The *indiscriminability view*, which says that what it is to hallucinate is to be in a mental situation that cannot be first-personally discriminated from a veridical perception.
- (ii) The *cognitive view*, which says that what it is to hallucinate an object is to be in a mental situation that lacks phenomenal character, but that produces in the subject the same cognitive effects that a veridical perception of an object would have produced in a rational subject with the same background mental situation.
- (iii) The *imaginative view*, which says that part of what it is to be a hallucination is to be an involuntary sensory imagining.

I argue in Appendix 6B that these accounts have a very hard time explaining the hallucinatory categoricity and hallucinatory intrinsicity data.

4. Some ways of conceiving of objects

Suppose that, while visiting the national parks in California, I see a redwood tree. I pace around it and take in the enormous trunk; I crane my neck and see the leaves high above. So far we have been focusing on how a perception such as this makes it possible to form a certain conception of *properties*. In addition to this, however, this perception makes it possible to form a multi-layered conception of *objects*.

Objects as enduring over time. For one thing, my perception positions me to conceive of an object (such as the tree) as something that *endures over time*. For a useful contrast, imagine a demon who creates a redwood-like object that exists for just a moment. The demon then destroys the object and creates a very similar one in the same spot, over and over, so quickly that these moments of creation and destruction cannot be seen. Or consider the *perdurantist* view that a redwood is merely a collection of distinct redwood-stages, each of which exists at a different time. My experience of the redwood tree positions me to form a conception of it that is incompatible with these possibilities – a conception of a single object that loses and gains properties over time. In fact, the datum is somewhat stronger than this: my perception of the redwood over time is veridical only if the redwood is a single object that endures over time. This does not entail that endurantism is correct, of course: perhaps my perception of the redwood is not veridical in this respect.¹⁰

¹⁰ I tentatively interpret Campbell (2002) as having this datum in mind. See Appendix 6D for further discussion.

Now that it is clear what the datum is, the explanation is simple enough. Suppose that it is possible to perceptually represent identity objects at various times; suppose that it is also possible to perceptually represent identity. Then it is possible for me to perceptually represent the redwood that I see now as identical to the redwood that I saw a moment earlier. This perceptual representation will be inaccurate if a demon has replaced the previous redwood with a new one, for the new redwood is not identical to the old one. The representation will also be inaccurate if a redwood is nothing more than a series of redwood-stages existing at different times, as the stages are not identical to one another. But wait – couldn't my perception represent, in perdurantist fashion, the self-identity of the entire *collection* of stages, rather than the identity of one stage with another? The pluralist can reasonably say that it could not: she can say that my perception at a given time cannot represent the entire collection of stages, which spans a great deal of time.

Objects as involving substrata. My perception positions me to form a conception of objects that has another aspect. When I see the redwood, I experience its shape, its coloration, and its many textures, and some metaphysicians would say that the tree is nothing more than a bundle of properties such as these. However, my experience positions me to conceive of the tree as more than that. It positions me to conceive of the tree as involving a *substratum* for these properties, i.e., a kind of entity that can *have* properties but that is not *itself* a property.

Note well: the claim is not that my perception positions me to conceive of the redwood as *being* a substratum. That cannot be right: I conceive of the tree as having some of its properties essentially, whereas a substratum does not have any of its properties essentially. The claim is rather that my perception positions me to conceive of the redwood as *involving* a substratum: I can conceive of a redwood as being a substratum *together with* some essential properties, where that substratum also happens to have some further accidental properties. Nor is the claim that

my perception is veridical only if the redwood involves a substratum. It is not. My perception could be perfectly veridical even if the tree were merely a bundle of properties. The claim is just that my perception *positions* me to conceive of the redwood as involving a substratum. (In this way, my conception of the tree as enduring over time is different from my conception of it as involving a substratum.)

These two aspects of my conception of the redwood should not be confused: they are independent. On the one hand, it is possible to conceive of a mere bundle of properties that persists over time. For example, it is possible to conceive of the tree as nothing more than a bundle of instances of brownness, cylindricality, and so on, while also taking these very property-instances to persist over time. On the other hand, it is possible to conceive of a series of momentarily existing but qualitatively similar substrata. Again, perhaps a demon could produce such a thing.

The present datum is also easy to explain, I think. The key point is that perception does not even seem to directly reveal substrata. When I see a redwood, my perception first-personally seems to confront me with that very object, as well as the various colors, textures, and so on that it has. But I do not seem to be confronted with a substratum; I do not seem to be directly confronted with something that has *none* of its properties essentially. Thus I suggest that the conception of a substratum is merely a theoretical posit formed on the basis of what we do directly perceive. Perception positions us to conceive of objects that can survive the loss of *some* properties. It is a small step to positing objects that can survive the loss of *all* properties – i.e., substrata.

Objects as mind-independent. There is some relationship between perception and the first-personal appearance of mind-independence. For instance, when I see the redwood, that object first-personally strikes me as *mind-independent*. Let us try to articulate this datum more precisely.

The datum is not that *whenever* a subject perceives an object, that object first-personally seems to be mind-independent. It is entirely possible to perceive an object that first-personally seems to be mind-*dependent*. The point can be illustrated with some of our examples from chapter 5. There we considered a subject who rubs her eyes and then opens them; she experiences phosphenes while also seeing some actual faint lights that are first-personally indiscriminable from phosphenes. This subject is seeing lights that are in fact mind-independent, but they first-personally seem to her to be mind-dependent. The same is true of the subjects who were told to imagine a blue banana and who then saw a faint projection of a blue banana on the wall. These subjects saw mind-independent images that first-personally struck them as mind-dependent.

At best, then, the datum is just that, *typically*, when a subject perceives an object, that object first-personally seems to be mind-independent. In fact, the datum is even weaker than this. When I see a redwood, the tree does not first-personally seem to be independent of *all* minds. For all that my experience bears on the matter, the existence of the tree might depend on the mind of other beings, such as God. It is just that the existence of the tree first-personally seems to be independent of *my* mind.¹¹

Here is the datum, then: in any typical perception, the existence of the perceived object first-personally seems not to depend on the mind of the subject herself.¹² How can the pluralist explain this datum?

My proposal is to extend the categorization account developed in chapter 5. The account was this:

¹¹ As Mackie (2019, §4.3) observes.

¹² To be precise, the datum is slightly broader. It applies to other targets of perception besides objects: property-instances, events, etc. But it is harmless to focus just on objects.

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.

I now add that, typically, when an experience is categorized as a perception, its target is simultaneously categorized as existing independently of the subject's mind. This concept of mind-independence (like the concept of perception) is a theoretical one that develops during the normal maturation of the subject. It is not hard to see how this might be made possible on the basis of perceptual learning. Imagine a toddler interacting with a tree that she sees. She moves around the tree, exploring it from different angles, moving closer and then further away. The pattern of sensory qualities that she experiences would best be explained by the hypothesis that she is seeing a tree that exists independently of her mind.

We can acquire some independent evidence for this approach by returning to a case that I mentioned in an appendix to chapter 5. 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 mind-independent tomatoes. You might take yourself to be having mere sensations – tomato-phosphenes, perhaps.

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.

Compare your experience of a tomato before the antics of the demon to your experience of a tomato afterwards. Both of these experiences are perceptions of mind-independent 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 present mind-independent tomatoes. Your second experience will seem to present mere tomato-phosphenes. What explains this difference?

The pluralist answers this question by appealing to the different histories of the subject. In particular, she suggests that perceptual learning is what teaches us which sensory gestalts are associated with mind-independent objects and which are not. Thus objects with precisely the same sensory qualities can be categorized as mind-independent or not, depending on what (apparent) perceptual learning the subject has undergone.

However, it has been argued that if perceptions were representational, then perceived objects could not first-personally appear to be mind-independent. One objection is that it is impossible to extract the conception of a mind-independent object from a mind-dependent representation; another objection is that if perceptions were representational, then they would

not be more fundamental than thought for grasping what an object is (or for grasping the mind-independence of an object). I respond to these objections in Appendix 6C.

Hallucination. Imagine an ordinary but unreflective undergraduate, Fiona. One day, a demon causes Fiona to have hallucinations of mundane life. Fiona hallucinates going to a philosophy class. She has never before considered whether objects endure (rather than, say, perduring); whether they involve substrata; and whether they are mind-independent. But she reflects on her hallucinatory experiences and comes to form such a conception of the objects that she seems to see around her. It seems obvious that this is possible: Fiona's hallucinations would equip her to form these conceptions just as well as matching perceptions (even if the hallucinations would not position her to have *knowledge* to the effect that the hallucinated objects are as she conceives them to be).

The pluralist can easily explain how this is possible. For she says that the subject forms various object-conceptions solely on the basis of having perceptual representations and deep awareness of certain kinds. These very same kinds of perceptual representation and deep awareness can occur in hallucination. The reader should immediately see how the details of this explanation will be filled in, and why this fact about Fiona's hallucination is problematic for the naïve realist.

And there is one last asymmetry. Fiona's hallucinations of a tomato will not just position her to conceive of roundness as a categorical and intrinsic property. It will position her to *know* that roundness is such a property. But it will not position her to know anything similar about *instances* of roundness, nor will it position her to know that "the" hallucinated tomato endures, that "it" involves a substratum, or that "it" is mind-independent. The pluralist explains this asymmetry by saying that Fiona is genuinely deeply aware of roundness even when she

hallucinates. However, Fiona is not perceptually related to any instances of roundness, nor to a tomato.

Conclusion. The pluralist can explain why perceptions, and even hallucinations, position us to conceive of objects in various distinctive ways.

7. Conclusion

The pluralist can explain why perception positions us to form a distinctive conception of certain *properties*: we can conceive of them as intrinsic and categorical because our deep awareness reveals this to us. The pluralist can also explain why perception positions us to form a distinctive conception of *objects*. Perception positions us to conceive of objects as enduring over time by representing the identity of objects experienced at different times. It seems to reveal to us objects that survive the loss of some properties, thereby positioning us to form a theoretical conception of objects that can survive the loss of all properties – i.e., substrata. And it positions us to conceive of objects as being mind-independent on the basis of perceptual learning. At the same time, pluralism has much more explanatory power than naïve realism: it can explain why certain *hallucinations* also position the subject to form these conceptions of properties and objects. All of this is to the credit of the pluralist theory.

Appendix 6A. Is it possible to represent objects as having categorical and intrinsic properties?

There is a significant representationalist component to the pluralist's account of our perceptual awareness of categoricity and intrinsicity. However, it has been argued that no purely representationalist theory can explain how perception makes us aware of the categorical and

intrinsic characters of objects, and there is a natural extension of the argument that targets pluralism. In this section, I respond to this extended argument.

Start by considering the version of the argument that targets pure representationalism. We can distinguish two possible types of perceptual representations of objects: *dispositional/extrinsic representations*, which are representations of objects as having merely dispositional and extrinsic properties, and *categorical/intrinsic representations*, which are representations of objects as having properties that are categorical and intrinsic as well. The pure representationalist wants to appeal to the latter. However, the thrust of the argument is that she cannot do so: the underlying functional facts cannot make it the case that perceptual representations are categorical/intrinsic rather than dispositional/extrinsic.¹³

The argument begins by assuming, for the sake of *reductio*, a pure representationalist account of object perception:

- C1. To perceive objects is simply to deploy perceptual representations of them.

¹³ This argument appears in Campbell (2002, pp. 145-153). The core of the argument is this (with numbering added to anticipate the way that I will regiment the argument shortly): “(C2) All we have, to differentiate between [dispositional and categorical representations], is what the subject makes of the representation. Does the subject interpret the representation as relating to the dispositional or as relating to the categorical? (C3) All that we have to appeal to, as constituting the subject's interpretation of the representation in one way or another, is the pattern of functional relations in which the representation stands to other representations and ultimately the actions of the subject. But it is very hard to see how the existence of such a pattern of functional connections could constitute the subject's interpreting the representation as relating to the categorical, rather than the dispositional, characteristics of the external stimuli. For the pattern of functional connections is determined only by the dispositional characteristics of the external stimuli. (C5) If you varied the intrinsic nature of the objects around the subject, but kept their dispositional characteristics constant, the very same pattern of functional connections would still be appropriate for the use of the subject's perceptual representations. And the actions of the subject would also be of the same type, in responding to and manipulating the affordances of the external stimuli. So (C6) on a Representationalist view, we seem to have no motive for going beyond [dispositional representations], to suppose that the subject is representing the categorical” (p. 150).

It is then claimed that:

- C2. Whether a representation of an object represents it as having categorical and intrinsic properties, or merely as having dispositional and extrinsic properties, is determined purely by how the subject *interprets* that representation.

The idea is that it is surely *possible* to represent an object as merely having dispositional and extrinsic properties. It is hard to see what else, besides the subject's interpretation, could *make it the case* that a representation is not of the object as having merely such properties, but rather of the underlying categorical and intrinsic properties.

The next premise is:

- C3. How the subject interprets a representation is determined purely by the pattern of functional relations between (i) the representation in question and (ii) the relevant inputs, outputs, and other mental states.

Again, the idea is that it is hard to identify anything else that could determine what the subject's interpretation is.

C2 and C3 entail:

- C4. Whether a representation of an object represents it as having categorical and intrinsic properties, or merely as having dispositional and extrinsic properties, is determined purely by the pattern of functional relations between (i) the

representation in question and (ii) the relevant inputs, outputs, and other mental states.

But notice:

- C5. The pattern of functional relations between (i) and (ii) is sensitive only to the dispositional and extrinsic properties that an object has; this pattern is insensitive to any further differences in the underlying categorical and intrinsic properties.

We can understand the idea here by comparing a normal subject seeing a tree (call her Nicky) with a subject who is merely a brain in a vat (call her Vicky). We may suppose that the brains of Nicky and Vicky are intrinsically identical at all times. But whereas the brain state underlying Nicky's perceptual experience is caused by an ordinary tree, the brain state underlying Vicky's perceptual experience (if this is indeed an experience) is caused by electrical impulses sent by wire from a computer.¹⁴

The thought is that Nicky and Vicky are functionally identical. That is, their brains are sensitive only to the extrinsic dispositions of the relevant external objects (the tree in Nicky's case; computers and wiring in Vicky's case). The fact that these objects have radically different categorical and intrinsic characters is irrelevant.

Now, C4 and C5 tell us that for any representations of external objects, their contents will be determined in a way that is insensitive to the categorical and intrinsic properties of those objects. And C1 tells us that to perceive is simply to deploy perceptual representations. Hence:

¹⁴ This is a mild elaboration of the case described in Campbell (2002, p. 149).

C6. To perceive objects is simply to represent them as having merely dispositional and extrinsic properties.

But C6 is plainly false! Nicky is plainly aware of the categorical and intrinsic character of the tree in front of her:

C7. To perceive objects is not simply to represent them as having merely dispositional and extrinsic properties.

C6 and C7 contradict each other. Thus we must reject our starting assumption, C1, which gives us:

C8. To perceive objects is not simply to deploy perceptual representations of them.

There ends the argument.

I accept the conclusion of this argument: as a pluralist, I think that object-perception involves not only perceptual representation, but also deep awareness of categoricity. But this does not yet get me off the hook. The problem is that I admit that the subject perceptually represents objects as categorical and intrinsic. But this view can be targeted by making minor modifications to the argument – in particular, the only claims changed below are C1, C6, C7, and C8:

C1*. In perception, the subject deploys perceptual object-representations. (Assumed for the sake of *reductio*.)

- C2*. Whether a representation of an object represents it as having categorical and intrinsic properties, or merely as having dispositional and extrinsic properties, is determined purely by how the subject *interprets* that representation.
- C3*. How the subject interprets a representation is determined purely by the pattern of functional relations between (i) the representation in question and (ii) the relevant inputs, outputs, and other mental states.
- C4*. Whether a representation of an object represents it as having categorical and intrinsic properties, or merely as having dispositional and extrinsic properties, is determined purely by the pattern of functional relations between (i) the representation in question and (ii) the relevant inputs, outputs, and other mental states. (Supported by C2-C3.)
- C5*. The pattern of functional relations between (i) and (ii) is sensitive only to the dispositional and extrinsic properties that an object has; this pattern is insensitive to any further differences in the underlying categorical and intrinsic properties.
- C6*. In perception, the subject does not deploy any perceptual object-representations that represent objects as having categorical and intrinsic properties; at best, her perceptual object-representations represent objects as having dispositional and extrinsic properties.

C1* is a commitment of pluralism, and C6* is supported more or less as before by C1*, C4*, and C5*. But, given the assumption of the form of representationalism expressed by C1*, C6* is supposed to be obviously false. What must be true, given C1*, is supposed to be:

C7*. In perception, the subject deploys perceptual object-representations that represent objects as having categorical and intrinsic properties.

C6* and C7* are contradictory. Thus we must reject the assumption that led us to this absurdity, C1*:

C7*. In perception, the subject does not deploy perceptual object-representations.

Now we have an argument whose conclusion really does target pluralism.

I respond that this argument misunderstands the pluralist theory in a subtle but important way. The pluralist thinks that C6* is true – she rejects C7* – thereby removing the contradiction. For she thinks that in perception, the subject’s perceptual object-representation does not, in and of itself, represent the object *either* as having categorical/intrinsic properties *or* as having dispositional/extrinsic properties. The object-representation does not represent the object *as* anything at all; it just *represents the object*, full stop.

The pluralist *does* think that in perception, the subject represents objects as having categorical and intrinsic properties. Her point is that this does not occur *because of object-representations alone*. Where do categoricity and intrinsicality enter the picture, then? The pluralist says that the subject is deeply aware of certain categorical and intrinsic properties. The categoricity and intrinsicality of these properties is thereby revealed to the subject. Though this awareness is entirely non-representational, the subject then puts it together with her object-representation to represent the object as having these categorical and intrinsic properties. To put the pluralist’s response in another way, then: to assert C7* is to assume that a perceptual

representation of objects as categorical must be fully grounded in what perceptual object-representations are like. The pluralist denies this assumption.

Even at this point, however, the concern is not fully allayed; the argument can be reformulated once more. The concern should now be that perceptual representations cannot represent *objects* at all; they can only represent the collections of dispositions associated with objects. This gives us the following version of the argument:

- C1**. In perception, the subject deploys a representation that represents either an object or a collection of dispositions. (Assumption for the sake of *reductio*.)
- C2**. Whether a representation represents an object or merely a collection of dispositions is determined purely by how the subject interprets that representation.
- C3**. How the subject interprets a representation is determined purely by the pattern of functional relations between (i) the representation in question and (ii) the relevant inputs, outputs, and other mental states.
- C4**. Whether a representation represents an object or merely a collection of dispositions is determined purely by the pattern of functional relations between (i) the representation in question and (ii) the relevant inputs, outputs, and other mental states. (Supported by C1**-C3**.)
- C5**. The pattern of functional relations between (i) and (ii) is sensitive only to what collections of dispositions an object has. It is insensitive to any further differences in which object has those dispositions.

C6**. In perception, the subject deploys a representation that does not represent objects; it represents only collections of dispositions. (Supported by C1**, C4**, and C5**.)

C1** is a commitment of pluralism, but C6** is inconsistent with pluralism. Thus, if pluralism is true, then pluralism is false. So:

C7**. Pluralism is false.

To my mind, this is the version of the argument that is the most threatening to pluralism. Still, I believe that the pluralist can offer a satisfying response. I am willing to concede everything up to C4**. What I resist is:

C5**. The pattern of functional relations between (i) and (ii) is sensitive only to what collections of dispositions an object has. It is insensitive to any further differences in which object has those dispositions.

Return to Nicky, our normal subject who sees a tree. I grant that, in principle, something radically different from a tree, such as a computer emitting electrical impulses, might have just the same narrow causal effects on any putative perceptual representations in Nicky's brain. From here, my opponent infers that the tree is not a *causal difference-maker* to Nicky's perceptual representations. But this inference is a poor one: one entity can be a causal difference-maker for another even if it is *possible* for some third entity to make the same difference. (Otherwise an entity could be a difference-maker only if its causal profile were *perfectly unique*.)

By analogy, suppose that a window was smashed because a baseball struck it, and the baseball struck it because Grace threw it. Grace's throwing the baseball is clearly a causal difference-maker for the window: had Grace not thrown the baseball, it would not have struck the window, and no window-smashing would have occurred. To be sure, the baseball could in principle have struck the window without Grace's help, for example by being swept up in a hurricane and then discharged in the general direction of the window. But that is beside the point! It is still Grace's baseball-throwing that *actually* caused the window-smashing.

Similarly, given the way that the world *actually* is, the tree *is* a causal difference-maker for Nicky's perceptual representations: had Nicky been looking at a radically different object, she would have had a radically different experience. Yes, it is *possible* for a radically different object, such as a piece of electrical wiring, to cause Nicky's brain to respond in the same way. But that is neither here nor there. It is the tree that *actually* causes certain neural processes to occur in Nicky. In virtue of this, Nicky's perceptual representation is a representation of the tree.

There might seem to be an important disanalogy between the cases of Grace and Nicky. In Grace's case, the difference-making in question is purely causal. By contrast, the pluralist understands Nicky's case as involving a causal difference *that then grounds a metaphysical difference*: the tree stands in a certain causal relation to Nicky's brain state, which in turn grounds the fact that Nicky perceptually represents the tree. However, the cases can be made analogous in this respect, as well: the fact that Grace causes the window to shatter might ground the fact that something bad has happened. Something bad would still have happened had the baseball been caught up in a hurricane and flung at the window. Nevertheless, Grace's throwing the baseball is in fact a difference-maker for the bad event.

I conclude that not only can pluralism explain the data about categoricity, but its explanation is much better than the standard naïve realist one.

Appendix 6B. Problems for some naïve realist theories of hallucination

Recall:

The hallucinatory categoricity datum: Take any perception that positions the subject to appreciate, in a special way, the categoricity of certain perceived properties. A matching hallucination will put the subject in the same position with respect to those same properties.

The hallucinatory intrinsicity datum: Take any perception that positions the subject to appreciate, in a special way, the intrinsicity of certain perceived properties. A matching hallucination will put the subject in the same position with respect to those same properties.

Can the naïve realist explain these data? The task is not an easy one, as we can see by briefly surveying some naïve realist views of hallucination.

One view, the *indiscriminability view*, is that what it is to hallucinate is to be in a mental situation that cannot be first-personally discriminated from a veridical perception.¹⁵ This view seems plainly inadequate: the *lack* of a discriminatory ability surely cannot explain a positive piece of understanding.¹⁶

¹⁵ Martin adopts a version of this view that is restricted to hallucinations that causally match veridical perceptions: “the disjunctivist insists that there is only a negative characterization of causally matching hallucinatory experience: it is nothing but a situation which could not be told apart from veridical perception ... this is then the fundamental mental character of the event” (2004, p. 72). Martin’s restricted view is equally subject to the criticism that I will press.

¹⁶ In a similar spirit, it has been argued that the present account of hallucination makes it impossible to explain other data. For example: that hallucinating red positions you to acquire

But wait – the indiscriminability theorist might reply that this lack of a discriminatory ability can position the subject to gain at least one piece of positive knowledge: knowledge that *either* she is perceiving a round object *or* it merely seems to her that this is so.¹⁷ The indiscriminability theorist might suggest that this positive knowledge is what puts Beatrice in a position to gain a substantive appreciation of the categoricity and intrinsicality of roundness.

I do not see how this last claim could be true. Suppose that I think that I have proved a particular mathematical result, *M*, but I am mistaken; my “proof” is spurious. Suppose further that I simply lack the acumen to identify, or even understand, where my “proof” has gone wrong – but *M* is true, and provably so. My inability to discriminate my “proof” from a genuine proof of *M* might position me to know that either I have proved *M* or it merely seems to me that I have proved *M*. Would this piece of positive knowledge position me to gain a substantive understanding of anything *mathematical* that I could not understand before? Evidently not. I do not see why things would be different in the case of perceptual experience.

But set that aside. The indiscriminability theorist faces a further problem. In order to know that *p*, the subject must *already* have the concepts required to entertain the proposition that *p*. Here is Beatrice: she is hallucinating a red round tomato. According to the indiscriminability theorist, that amounts to her being in a mental situation that cannot be first-personally discriminated from a veridical perception of a red round tomato. Grant for the sake of argument that, *if Beatrice had the requisite concepts*, being in this mental situation would position her to know that (i) either she is perceiving a round object or it merely seems to her that this is so. Grant also that (ii) if she were to know this, then she would be able to appreciate, in a special way, the

knowledge about redness (Johnston 2004) or at least to form beliefs about redness (Pautz 2007, p. 525), and that hallucinating a butterfly positions you to know that you are not veridically perceiving a sausage (Siegel 2008, pp. 211-214). I am tentatively sympathetic to these arguments.

¹⁷ This style of response is recommended in Soteriou (2016, p. 180).

categoricity and intrinsicity of roundness. Still, Beatrice does not yet have the concept of roundness, so she is unable to come to know (i), or even to believe (i).¹⁸ Thus she is unable to satisfy the antecedent of (ii).

From this, I conclude that the naïve realist who adopts the indiscriminability view cannot explain the hallucinatory categoricity and hallucinatory intrinsicity data.

But there are other views about hallucination that the naïve realist might adopt instead. One such view is that what it is to hallucinate an object is to be in a mental situation that lacks phenomenal character, but that produces in the subject the same cognitive effects that a veridical perception of an object would have produced in a rational subject with the same background mental situation. The idea is that what makes a mental situation a hallucination of a maple tree is the fact that it produces in the subject (e.g.) the belief that she is seeing the pointed leaves, broad branches, and stout trunk of a maple tree. Call this the *cognitive view* of hallucination.¹⁹

Beatrice is hallucinating a red round tomato. On the cognitive view, this amounts to her enjoying a mental state or event that produces in her the cognitive effects that would have been produced in a rational subject seeing a red round tomato. Precisely which effects are those? One effect would be the belief that she is seeing a fruit, one that would taste a particular way. But it is of no use for the cognitive theorist to appeal to this effect – such a belief obviously cannot, by itself, position the subject to appreciate the categoricity and intrinsicity of roundness in a special way.

Of course, there is a further effect that seeing a round tomato would have on a rational subject: it would thereby position her to appreciate, in a special way, the categoricity and intrinsicity of roundness. But the cognitive theorist cannot appeal to *this* effect – the very

¹⁸ She has, as epistemologists like to say, *propositional* justification but not *doxastic* justification for believing (i).

¹⁹ This account appears in Fish (2009, pp. 94–95).

question at hand is what puts a hallucinating subject in such a position. The cognitive theorist claims that *if* a subject is in a position to gain this appreciation (among other things), *then* she is hallucinating. But that does not tell us how it is possible to end up in such a position to begin with.

Here is another way to put the problem. We can all agree that for a subject in a certain background mental situation, a neural state that does not partly constitute a perception might *just happen* to have in her many of the effects that a perception would have had. The cognitive theorist says that any such neural state is a hallucination. Fine! But if one effect of a neural state is the ability to appreciate the categoricity and intrinsicity of roundness, then the cognitive theorist still owes us an account of how a neural state could have *that* effect.

The cognitive theorist might try an alternative strategy: she might deny that having a hallucination puts the subject in a position to appreciate the categoricity and intrinsicity of roundness. Instead, the cognitive theorist might say, such a subject merely *first-personally seems* to be in such a position.

This response denies what I take to be an obvious truth. Imagine that you have never seen anything round before. Then you hallucinate a round tomato. It strikes me as undeniable that this puts you in a position to appreciate the categoricity and intrinsicity of roundness. The point is driven home when we consider what you will say even if you understand at the time that you are merely hallucinating: you will say that, yes, there is no particular round object that you are experiencing, but still you now appreciate in a new way certain facts about the property of roundness itself.

Let us turn to a third naïve realist view of hallucination, then. According to the *imaginative view*, part of what it is to be a hallucination is to be an involuntary sensory imagining. This is not

intended as a full account of the nature of hallucination: some involuntary sensory imaginings are not hallucinations.²⁰

Perhaps the partial account is enough to explain the data: my opponent need only add that episodes of sensory imagination can position the subject to appreciate the categoricity and intrinsicity of certain properties. On this view, the idea would presumably be that what a categorical ground is. But this just pushes back the problem: it now remains to be explained how episodes of sensory imagination can do this.

The imaginative theorist might respond that sensory imaginings are representations *of perceptions*.²¹ As a naïve realist, the imaginative theorist might think that she already has an account of how perceptions position us to appreciate, in a special way, the categoricity and intrinsicity of properties such as roundness. She may add that in virtue of being representations of perceptions, sensory imaginings can do the same work.

I see several problems here. First, I have already identified several weaknesses in the naïve realist account of how *perceptions* put us in this position, at least relative to the pluralist's account. The imaginative theorist inherits these relative weaknesses. But set this aside for the sake of argument. The problem remains that the subject of our example, Beatrice, has never seen anything round; she has never stood in the naïve relation to an instance of roundness. It is unclear, then, how it could be possible for her to represent a perception *of roundness* at all. But we are currently entertaining the view that a hallucination of roundness is a sensory imagining –

²⁰ The imaginative view of hallucination is defended carefully in Allen (2015). On pp. 295-296, Allen mentions some examples of involuntary sensory imaginings that are not hallucinations. However, he does not make any attempt to complete the account by offering a full characterization of the nature of hallucination.

²¹ This view of sensory imagining is defended in Martin (2001, pp. 270-273) and (2002, p. 404 and p. 407). Allen entertains this view without committing to it (2015, p. 289).

i.e., a representation of a perception – of roundness. Thus, the imaginative theorist should now predict that Beatrice cannot hallucinate roundness. This is evidently incorrect.

In sum, none of these naïve realist accounts of hallucination – the indiscriminability view, the cognitive view, or the imaginative view – can adequately explain the hallucinatory categoricity and hallucinatory intrinsicity data.

Appendix 6C. Objections involving our appreciation of mind-independence

Some have worried that a representationalist cannot possibly explain why perceived objects typically first-personally seem to be mind-independent. This core concern has been developed in several different ways. Let us consider whether these might threaten pluralism.

One objection is that on the representationalist account, perception provides the subject with nothing more than a perceptual representation, i.e., a mind-dependent image. But the subject cannot extract a conception of mind-independence from something that is mind-dependent. Thus, representationalists cannot explain why perceived objects typically first-personally seem to be mind-independent.²²

This objection fails to undermine even a standard representationalist account. We can appreciate the point by identifying the precise sense in which, according to the standard

²² As Campbell puts it, “On the [representationalist theory], all that experience of the object provides you with is a conscious image of the object – the image which bears the representational content. The existence of that conscious image is in principle independent of the existence of the external object. The existence of the image, though, is dependent on the existence of the subject who has the conscious image. So if your conception of the object was provided by your experience of the object, you would presumably end by concluding that the object would not have existed had you not existed, and that the object exists only when you are experiencing it. We cannot extract the conception of a mind-independent world from a mind-dependent image” (2002, p. 121).

representationalist, perception *provides* the subject with nothing more than a perceptual representation. The claim is just that to perceive is, in part, to perceptually represent. This is *not* the claim that the subject is *perceptually aware* of nothing more than a perceptual representation. To begin with, on standard representationalist accounts, the subject need not be perceptually aware *of* perceptual representations at all.²³ More importantly, the standard representationalist holds that the subject is directly perceptually aware of – that is, she directly perceptually represents – *particular mind-independent objects and their property-instances*. The pluralist agrees. So the problem of deriving a conception of mind-independent entities from awareness of mind-dependent entities does not even arise.

Now consider a pair of related arguments. This first of these arguments runs as follows (*MI* stands for *mind-independence*):

- MI1. If perceiving objects is merely a matter of representing them, then perception is no more fundamental than other forms of thought for grasping what an object is.
- MI2. Perception is more fundamental than other forms of thought for grasping what an object is.
- MI3. Perceiving objects is not merely a matter of representing them.²⁴

²³ Although the *higher-order representationalist theory* holds that the subject perceptually represents *both* mind-independent entities *and* her own perceptual representations. See [cite].

²⁴ As Campbell puts it, “Experience is what explains our grasp of the concepts of objects. But if you think of experience as intentional, as merely one among many ways of grasping thoughts, you cannot allow it this explanatory role At this point we do not have any way of explaining why there should be anything fundamental to our grasp of concepts about experience of objects.... And at this point the question whether something essentially mind-dependent could provide for the conception of a mind-independent world really does seem forceful” (2002, p. 122).

The support for MI1 is that there are many forms of representation, from speech to facial expressions to signs, and the pure representationalist treats perception as merely one more form of representation.²⁵ MI2, meanwhile, is intended to be obviously true. It is one thing to form a demonstrative concept of a particular rabbit just by seeing its tracks in the garden. It is quite another to form such a concept on the basis of actually seeing the rabbit. Only the latter positions you to appreciate, in an especially direct way, what a rabbit.

The pluralist has nothing to fear from this argument, for she accepts its conclusion. She says that perceiving an object requires not only representing it, but also having deep awareness that reveals part of the essence of certain qualities of the object. The latter is not found in ordinary thought, and it is what positions us to know *what the object is*, in a certain special way – that was the lesson of my discussion of strong singular character in chapter 4.

Here is a similar argument:

MI1*. If the first-personal perceptual appearance of objects as mind-independent is merely a matter of perceptually representing (or categorizing) those objects as mind-independent, then perception is no more fundamental than other forms of thought for appreciating the mind-independence of an object.

MI2*. Perception is more fundamental than other forms of thought for appreciating the mind-independence of an object.

²⁵ As Campbell puts it, “To see an object is, on this conception, to grasp a demonstrative proposition. There are many ways in which you can grasp a proposition: you can grasp it as the content of speech or as the meaning of a wink or a sigh. One way in which you can grasp a proposition is as the content of vision” (2002, p. 121).

MI3*. The first-personal perceptual appearance of objects as mind-independent is not merely a matter of perceptually representing (or categorizing) those objects as mind-independent.²⁶

The premises can be supported much as before. Regarding MI1*, the idea is again that perceptual representation and categorization are each just further forms of representation, on a par with thought, speech, and signs. And MI2* is meant to be obviously true.

My response to this argument depends on how the phrase *appreciating the mind-independence of an object* is to be understood.

Perhaps the idea is just that seeing a tomato positions the subject to know *which* object she is taking to be mind-independent. On this interpretation, what is special about perception is the way that it relates the subject to the object, not to the property of mind-independence. In that case, I would reject premise MI1*. I have already explained how the pluralist can explain this special feature of perception: this is a matter of the subject's having partly essence-revealing awareness of the object's properties.

But there is a bolder interpretation, on which seeing a tomato positions the subject to appreciate the mind-independence of objects in a much stronger sense. The idea is that seeing a tomato *presents* the subject not only with the tomato, but also with its mind-independence. If this is the idea, then I reject MI2*. Seeing a tomato presents the subject with the tomato and its instances of red*, but it does not even seem to present her with the mind-independence of the

²⁶ This is an alternative way of interpreting Campbell's remarks in the footnotes above. This interpretation is driven by the fact that, at times, Campbell seems to emphasize not just our grasp of objects themselves, but of their mind-independence. For example, at one point Campbell says, "The objection to the common factor view is that on it, experience of objects could not be what explains our having the conception of objects as mind-independent" (2002, p. 121). (Here Campbell is stating an objection that he takes from McDowell, but the argument from the footnotes above is intended to develop McDowell's objection.)

tomato. Here is a way of driving the point home. Recall, the example of the subject who begins to mistake her tomato-perceptions for tomato-phosphenes, and compare this subject to an ordinary subject who sees a tomato and correctly takes it to be mind-independent. Does the experience of the latter subject present her with an additional element, *mind-independence*? Surely not.

I conclude that these objections are unsuccessful.

Appendix 6D. Campbell on the conceptions of properties and objects made possible by experience

I have said that experience positions us to form certain conceptions of properties and of objects. These claims are inspired by ideas from John Campbell's book *Reference and Consciousness* (2002). However, Campbell's discussion betrays some serious confusions. Here I attempt to develop a charitable interpretation of what he has in mind.

One problem is that Campbell speaks over and over of categorical *objects* (pp. 138-145). But only *properties* can be categorical. What does he mean, then? There are several clues.

One is that Campbell thinks that the categoricity of objects explains why they tend to transmit marks over time:

[First, a]n object at a later time is capable of bearing marks transmitted by the way it was earlier.... Suppose that while at school you carve your initials on a desk. When you revisit years later, there they still are. In this case, the identity of the object over time seems to be the categorical ground of its potentiality to sustain this kind of marking interaction. The reason why the marks are there is that it's the same desk.... We need the notion of a categorical object, whose categorical identity can be the ground of its complex of dispositions to interact with other objects. (p. 140)

A second clue is that Campbell thinks that the categoricity of objects explains why they tend to have a range of correlated effects at a given time:

The second type of phenomenon in explaining which we use the identity of macroscopic objects has to do with the tendency of macroscopic objects to produce ranges of correlated effects. For example, one person addressing a group of people may produce a range of effects in them: there may be a shared sense of lassitude, or on the other hand intense intellectual excitement, after such an encounter. What explains the correlation in the effects produced is, in part, that it was one and the same person addressing all these people. (p. 140)

A third clue is that Campbell consistently elides the distinction between the intrinsic and the categorical. For example, he holds that “[e]xperience is experience of the categorical,” but he calls this the *Intrinsicness Condition* (2002, p. 137). Similarly, he says this:

We would ordinarily regard shape properties as the paradigmatic *categorical* properties of objects. Roundness is the reason why the thing tends to roll if suitably propelled and so on; it is the *intrinsic* ground of the complex functional state described. (p. 139, emphasis added)

I propose the following interpretation: Campbell’s most fundamental thought is that perception positions us to think of objects as having certain *intrinsic properties*. This would explain why Campbell focuses on the tendencies of objects to transmit marks over time and to have a range of correlated effects at a given time, for the intrinsic properties of an object can frequently explain these facts. At the same time, many of the intrinsic properties of objects – perhaps all of their most explanatorily fundamental ones – are categorical. Thus, this interpretation would also explain why Campbell might run together intrinsicity and categoricity.

But there is a further interpretive problem: Campbell repeatedly says that perception positions us to conceive of the *identity* of an object²⁷, but it is unclear what he means by this. Everything is identical to itself, so what is the difference between being confronted with the identity of an object and being confronted with the object itself? Perhaps Campbell (sometimes) has in mind the identity of an object across time – that is, perhaps he means that perception positions us to conceive of objects as persisting over time.

Yet this reading runs into problems. For Campbell seems to think of the categoricity and identity of an object as one and the same: “How is it that we have this conception of the identity of a macroscopic object, as the categorical ground of these dispositions ...?” (p. 141). But whether objects persist over time is orthogonal to whether they have categorical properties. Return to the example of the demon who repeatedly produces a redwood-like object and immediately destroys it, replacing it with something very intrinsically similar to it. Every redwood-like object in this series has categorical properties, but none of these objects persists over time.

Perhaps this problem can also be solved if we interpret Campbell as focusing, most fundamentally, on the intrinsic character of objects. For we can then read him as thinking that *haecceities* – which distinguish qualitatively identical entities – are intrinsic. In that case we can read Campbell as emphasizing that seeing a redwood over time reveals another intrinsic aspect of that object: the fact that it is the *same* object that persists over time.

²⁷ Again, see Campbell (2002, pp. 138-145).