

Precisifying 'Perception':

In Defense of a Pragmatic Stance on the Metaphysics of Perception

What is perception? More specifically, what can be said about the relation a subject stands in when they are perceptually related to some particular? Particularists hold that a subject's perceptual state brought about by being perceptually related to some particular is (in part) constituted by that particular, whereas generalists claim that all such perceptual states are wholly constituted by general elements. Here, I compare Schellenberg's (2018) argument for particularism against Lewis' (1980) and Harman's (1990) arguments for generalism. The debate, I argue, is a paradigmatic instance of what Plunkett (2015) calls 'metalinguistic negotiation'; the interlocutors use (rather than mention) the word 'perception' in conflicting ways to express normative views about how that term should be used. The upshot is that if the debate between particularism and generalism is a metalinguistic negotiation, then there is good reason to adopt a pragmatic stance about the nature of the connection between our perceptual states and the objects they are directed towards.

In part one, I summarize Schellenberg's main argument for the particularity thesis by considering what it is to be perceptually related to a white cup. The roles of perception in fixing reference on the cup, allowing for singular thought about the cup, and determining accuracy conditions of the perception of the cup, can all be explained by the thesis that perception necessarily involves discriminating and picking out particulars. Furthermore, Schellenberg argues, if perception necessarily involves discriminating and picking out particulars, then our perceptual states must be (partly) constituted by those particulars.

In part two, I argue that (like particularism) generalism can also account for the relevant roles of perception by positing a causal relation between the particulars perceived and the general features making up our perceptual states. The serious problems facing a generalist solution of this sort arise only if we uncharitably assume that generalists hold perceptual states to be wholly constituted by abstract non-spatiotemporally located properties. However, if—as I will argue—generalism merely claims that perceptual states are wholly constituted by mental states, then the only substantial dispute between particularism and generalism is whether the mind-independent particulars which are causally responsible for our perceptual states are, in fact, constitutive of those perceptual states. As we will come to see, the antecedent of the latter conditional is supported by a charitable reading of Lewis (1980) and Harman (1990), both of whom Schellenberg cites as canonical generalists.

Finally, in part three, I argue that if we take the dispute between particularism and generalism to be about whether the particulars which are causally responsible for our perceptual states are constitutive of those perceptual states, then the dispute is not about an ‘object-level issue’, but rather, a ‘representational-level issue.’ Hence, the issue is not whether perceptual states are, *by nature*, constituted by the particulars they are related to, but rather, whether we should define ‘perceptual states’ in terms of the particulars they are related to or in terms of purely general features. Following Plunket’s (2015) characterization of metalinguistic negotiations, the latter conclusion is supported by the fact that 1) the interlocutors are engaged in an actual disagreement, and 2) there is good reason to believe each interlocutor consistently means different things when they use the term ‘perceptual state’.

Part One

The particularity thesis (PT), as stated by Schellberg (2018), is the thesis that:

(PT) A subject’s perceptual state M brought about by being perceptually related to the particular α is constituted by α . (14)

Most arguments for (PT) have focused on the role that perception plays in our cognitive and epistemic lives—grounding singular thoughts, determining accuracy conditions, etc.—arguing that (PT) follows by an inductive inference to the best explanation. Going a step deeper, Schellenberg purports to explain *how* perception plays the latter roles, arguing that “perception is constitutively a matter of employing perceptual capacities that function to discriminate and single out particulars” (pg. 29). Here, I briefly explain each of these arguments and potential generalist responses, arguing that generalists can also provide a constitutive definition of perception which adequately explains the aforementioned roles.

1.1 Grounding Singular Thoughts

The argument from singular thoughts starts with the premise that perception gives rise to singular thoughts (i.e. thoughts about particular objects). For instance, the thought that “that teacup is white”—which I form on the basis of looking at the teacup sitting on the desk next to me—is plausibly about *the* particular teacup on my desk. Thus, the argument proceeds, my perception of the teacup is partly constituted by the teacup.

A slightly simplified version of the argument is as follows:

1. My perceptual state, P, brought about by perceiving the teacup on my desk gives rise to my singular thought about the teacup on my desk.
2. If P gives rise to my singular thought about the teacup on my desk, then P is partly constituted by the teacup on my desk.
3. Therefore, my perceptual state, P, is partly constituted by the teacup on my desk.

The teacup is for illustrative purposes, but this argument can be generalized to any perceptual state, yielding the particularist's desired conclusion.

It is reasonable to assume that premise one is true, as it is not implausible that being perceptually related to objects gives rise to 'singular thoughts' about those objects. Hence, premise two is the more controversial premise, and all arguments for it are inferences to the best explanation (e.g. my ability to have a singular thought about x is best explained by my perceptual state's being partly constituted by x). But, as Schellenberg notes herself, these arguments to the best explanation are weak, considering that the generalist can offer a compelling explanation by 1) restating their claim that perceptual states are constituted only by general elements and 2) accounting for singular thoughts by positing a causal relation between those general elements and the perceived particular.

Responding to the latter objection, particularists might argue that accounting for singular thoughts by positing a causal relation between general elements of the perceptual state and the perceived particular involves a strange discontinuity. Schellenberg (2018) writes,

It is odd to posit that perceptual content is general, only to maintain that the content of beliefs based on perception can be singular via the very causal relations to particulars operating in perception itself. After all, if the causal relations to perceived particulars can make the content of perceptual beliefs singular, why would they not make the content of perception itself singular? (pg. 21)

But this line of argument poses no serious problem for generalists, who can respond in one of the following two ways. First, generalists can simply reject that the discontinuity between the content of perception and thought is a serious problem. Singular thoughts—which are brought about by a causal relationship with some perceived particular—have singular content in virtue of the fact that the person causally related to the particular takes themselves to be causally related to that particular, though their perceptual state is not in fact constituted by that particular. So, for instance, the content of “that teacup is white” would be something like ‘*whatever particular is causing my current perceptual state [which is wholly constituted by general elements] is white.*’ Alternatively, the second response available to the generalist is to deny that there is a discontinuity between the supposed singular content of perception and thought. For instance, the generalist could argue that both perceptual states and the thoughts they give rise to have singular content, but that to have singular content is not to be constituted by particulars. On this view, for a perceptual state to have singular content about a particular is not for it to be *constituted* by that particular, but rather, merely *caused* by it.

In sum, the argument from singular thought fails to provide compelling evidence for particularism, relying on a weak intuition that particularism is the best explanation of the human ability to have singular thoughts. As shown above, the generalist's explanation of singular thought is equally (if not more) compelling than the particularist's. Similar arguments for particularism give inferences to the best explanation of other proposed roles of perception, but they each face an analogous weakness; namely, the weakness is that particularism is not clearly the best explanation. Generalists offer a compelling explanation by positing a causal relation to perceived particulars.

1.2 Determining Accuracy Conditions

The argument from accuracy conditions starts from the premise that when we perceive a particular, the conditions under which that perception is accurate are determined by the perceived particular. Thus, the argument proceeds, our perceptual states are constituted by the particulars that determine their accuracy conditions. This argument can be elucidated by returning to the teacup example:

1. My perceptual state, P, brought about by perceiving the teacup on my desk, has accuracy conditions that are determined by the particular teacup on my desk.
2. If P has accuracy conditions that are determined by the teacup on my desk, then P must be partly constituted by the teacup on my desk.
3. Therefore, my perceptual state, P, is partly constituted by the teacup on my desk.

The intuition behind this argument is that the accuracy conditions of my perception of the teacup (i.e. the conditions the world must satisfy for my perception of the teacup to be accurate) specify not just that there is some teacup on my desk, but that there is a particular teacup on my desk. For the perception to be accurate, the specific teacup tracked by my perceptual state must be on the desk.

Of course, the generalist can simply respond by denying that the accuracy conditions of our perceptual states specify exactly which particular must be present. For instance, my perception of the white teacup is accurate iff there is some white teacup there, but exactly which teacup is present is irrelevant to the accuracy of the perception. But the particularist has a compelling rebuttal. Schellenberg writes,

Suppose I look at a slanted mirror and have a visual experience that is caused by a white cup on my left reflected in the mirror. Not knowing that I am looking into a mirror, it seems to me that the cup is straight ahead. As it happens, there is a white cup behind the mirror just where it seems to me that there is a cup. Intuitively, the content of my perceptual state is inaccurate. (pg. 24)

So, contrary to generalism, this rebuttal states that specifying the accuracy conditions of perceptual states in terms of purely general features is insufficient. Driving this rebuttal is the intuition that perceptual states can be inaccurate even if there is some object which has the features they represent, as in the above case of the reflected teacup. But still, the generalist need not be swayed; they can simply reject the intuition that the perception of the reflected cup is inaccurate.

1.3 The Particularity Argument

Both of the arguments above—the argument from singular thought and the argument from accuracy conditions—employ an inductive inference to the best explanation to establish the truth of particularism. But the inferences fail in both cases because generalists also provide plausible explanations of the explanandum. Recognizing this shortcoming, Schellenberg puts forth a novel *particularity argument*, glossed roughly as follows:

1. When I perceive the teacup on my desk, I discriminate and single out the teacup on my desk (as a consequence of being perceptually related to the teacup.)
2. If I discriminate and single out the teacup on my desk, then my perceptual state P brought about by perceiving the teacup on my desk is constituted by discriminating and singling out the teacup on my desk.
3. If P is constituted by discriminating and singling out the teacup on my desk, then P is (partly) constituted by the teacup on my desk.
4. Therefore, my perceptual state P brought about by perceiving the teacup on my desk is constituted by the teacup on my desk.

(adapted from pg. 24-25)

This argument aims to avoid the faulty inductive inferences of the first two arguments present above, instead opting to argue that perceptual states are, by nature, constituted by the particulars which they are causally related to.

To evaluate the particularity argument, first note that premise 1 is fairly plausible. Defending premise 1, Shellenberg writes, “it is unclear what it would be to perceive a particular without at the very least discriminating and singling it out from its surround” (pg. 25). Indeed, perception seems to at the very least involve discriminating and singling out particular features of our environment, for what could it mean to perceive something if not to at least identify it as separate from other objects? Though it is possible to deny this, let us accept premise 1 to give Shellenberg the benefit of the doubt.

Premise 2, on the other hand, is not uncontroversial and can be responded to in a couple different ways. First, the generalist can deny 2 by arguing that the perceptual state P

brought by perceiving the teacup is not constituted by the teacup, but rather, merely causally related to it. This brings us back to a similar problem faced by the arguments above—generalism offers a perfectly plausible alternative explanation. Schellenberg resists this by stating that this explanation is “ad hoc and has no bearing on perceptual discrimination” (pg. 28), but this response could equally apply to particularism. Just as the denial of premise 2 seems to beg the question, so too does its affirmation.

Alternatively, Schellenberg argues, the generalist could respond by arguing as follows:

[...] we do not discriminate between property-instances, but rather between properties. Properties are not particulars, but rather universals. [Since] we discriminate between properties, there is no need to say that perceptual states are constituted by particulars. (pg. 27)

However, if the generalist pursues this strategy, then they commit themselves to the view that perceptual states are composed of abstract universal, or in other words, non-spatiotemporal and noncausal properties. But, as Schellenberg is apt to point out, “It is not clear what it would mean to perceptually discriminate between entities that are neither spatiotemporally located nor causally efficacious” (pg. 27). To discriminate and single out a property is to pick out something at a specific point in space and time, and it is unclear how we could pick out something at a specific point in space and time with platonic properties which do not occupy a position in space and time. This generalist response fails.

To recap, the argument from singular thought (presented in 1.1) and the argument from accuracy conditions (presented in 1.2)—both of which rely on an inference to the best explanation—fail to provide strong evidence for particularism. The problem is that, in each case, the generalist has equally if not more compelling explanation for the explanandum. On the other hand, the particularity argument does not rely on an explicit inference to the best explanation, instead arguing that particularism follows from the fact that perception, by nature, involves the process of discriminating and singling out particulars. The first premise of the argument is plausible and relatively uncontroversial, but the second premise is contentious. Lastly, the third premise, which has yet to be discussed, is defended largely on the basis that “To think otherwise would be to sever the link between the function of the capacity and its output” (pg. 26). For present purposes it should suffice to say that the generalist has several strong objections to this defense, first being the claim that capacities do not have to be defined in terms of their outputs, and second being the claim that even if capacities must, for some reason, be defined in terms of their outputs, it still does not follow that those capacities are comprised of their outputs.

The upshot is simply that only one of the arguments considered above gives any strong reason to prefer generalism to particularism. In all cases, generalism seems to entail (and hence, explain) all of the same phenomena as particularism. The only strong argument against generalism (and for particularism) has been the particularists response to the generalist's second objection to premise 2 of the particularity argument. If generalists assume that perceptual states are composed of abstract non-spatiotemporally located properties, then generalism cannot explain how we discriminate and single out objects. But this feels like a strawman. In the next section I consider whether the people Schellenberg cites as generalists really do make radical claims about properties that Schellenberg wants to reject. If so, then perhaps there is good reason to prefer particularism to generalism. But if not, then perhaps the argument is, as I argue, a paradigmatic metalinguistic negotiation.

Part Two

So, Schellenberg's strongest objection to generalism rests on the faulty assumption that the view assumes perceptual states to be constituted by non-spatiotemporally located properties, rather than property instances. Loosely speaking, this is like assuming that all generalists are Platonists who believe that our perceptual states are composed of abstract properties which participate in some general form. (Hence the reason Schellenberg explicitly endorses an Aristotelian view on which "properties are understood in terms of their instances") (pg. 15). But to characterize the debate between generalism and particularism (roughly) in terms of differences between Platonic and Aristotelean metaphysics of properties is to fundamentally miss the point.

In this section, I draw from Harman (1990) and Lewis (1980)—both of whom Schellenberg cites as canonical generalists—to show that Schellenberg's characterization of generalism is uncharitable. A more accurate characterization, I argue, is that particularists take the objects of our perceptual states (e.g. mind-independent events, universals, property instances, etc.) to be constitutive of those states, whereas generalists deny this, holding that our perceptual states are wholly constituted by our mental states. Hence, generalism is 'general' not because it assumes that our perceptual states are constituted by general non-spatiotemporally located properties, but rather, because perceptual states are more general than the mind-independent objects which cause them (e.g. cup_1 is distinct from cup_2 , yet both cause the same perceptual state).

2.1 Harman's 'Generalism'

If Schellenberg's characterization of generalism is charitable and accurate, then the people she cites as canonical generalists must make the claims that Schellenberg attributes to generalism. Here, I focus on Harman (1990), arguing that he does not, in fact,

make the claims Shellenberg attributes to generalism, and hence that Schellenberg's characterization of generalism is uncharitable and inaccurate.

Harman does not make any claims pertaining to whether perceptual states are (or are not) wholly constituted by non-spatiotemporally located properties, nor does he make any explicit reference to generalism or particularism. Rather, the paper is a defense of functionalism—a theory of mind which defines mental representations in terms of their functional role—against various objections which claim that we are directly aware of intrinsic features of our experience and that there is no way to account for this on a functionalist view. Ultimately, Harman responds to these objections by arguing that they result from a failure to carefully distinguish between properties of the objects of experiences and properties of the experiences of those objects.

To understand Harman's view and how it relates to generalism, consider the following argument against functionalism:

1. If functionalism is true, then all qualities of experience can be defined in terms of relations between perceptual inputs, mental states, and behaviors.
2. But when you attend to a pain, you attend to an intrinsic quality of experience (i.e. a quality of the experience in itself, and not a relation between things). Hence, there exists a quality of experience which cannot be defined in terms of relations between perceptual inputs, mental states, and behaviors.
3. Therefore, functionalism is false.

Harman accepts the first premise as an accurate characterization of functionalism, but rejects the second premise, thereby defending functionalism. Premise two is false because when you attend to a pain, you are not attending to an intrinsic quality of experience, but rather, a property of an object of experience. Harman (1990) writes,

[The] happening in your leg that you are presented with is the intentional object of your experience; it is not the experience itself. The content of your experience is that there is a disturbance of a certain specific sort in your right leg. The intentional object of the experience is an event located in your right leg. The experience itself is not located in your right leg. If the experience is anywhere specific, it is somewhere in your brain. (pg. 40)

So, the second premise is false because the experience of pain does not have an 'intrinsic' quality that is directly accessible. Rather, the supposed 'intrinsic' aspect of pain is not a property of the experience of pain, but rather, the intentional object of the experience—namely, the event happening in your leg.

To elucidate these confusing distinctions between properties of experiences, properties of the contents of experience, and properties of the intentional objects of experiences, Harman introduces a helpful analogy between paintings and experiences. Consider a painting of a unicorn. The unicorn painting is analogous to an experience, and what the painting represents is analogous to the content of the experience. The painting represents a unicorn, and so the 'content' of the painting is a unicorn. But unicorns do not exist, and so the content of the painting does not pick out an intentional object. Harman writes,

The unicorn is pictured as having four legs and a single horn. The painting of the unicorn does not have four legs and a single horn. The painting is flat and covered with paint. (pg. 35)

Analogously, the experience a person has when they attend to a pain in their leg does not have any intrinsic qualities that the person is directly aware of; the experience is an occurrence in the person's brain. Rather, what the person is aware of is an occurrence in their leg which is represented by the content of their experience.

A notable consequence of the latter characterization of pain is that, analogous to the fact that a painting of a unicorn can represent—and hence, possess the intentional content of—a nonexistent unicorn, it is possible for an experience of pain in your leg to have the intentional content that there is a disturbance in your leg without there actually being such a disturbance. This provides an interesting explanation of the relationship between normal perceptual experiences on the one hand, and veridical hallucinations and phenomena like phantom pain on the other. In the latter cases, people's experiences have contents which fail to pick out any intentional object.

As evidenced by the analysis above, Harman is committed neither to the idea that perceptual states are constituted by non-spatiotemporally located properties nor property instances. In fact, the question about whether perceptual states are constituted by non-spatiotemporally located properties seems irrelevant to Harman's argument, which is focused on defending a functionalist view of the mind by emphasizing the crucial distinction between properties of experiences and properties of the intentional objects of experiences. So, though Schellenberg cites Harman as a canonical generalist, Harman does not make the claims that Schellenberg attributes to generalism. Again, in response to this discrepancy, I suggest that what actually characterizes generalism is the claim that perceptual states are wholly constituted by mental properties. Hence, according to the generalist, so long as the mental properties you have remain constant, then you remain in the same perceptual state regardless of whether you are 1) causally related to a particular x, 2) causally related to a distinct particular y, or 3) not causally related to any particular.

On the latter definition, generalism follows intuitively from functionalism. Nevertheless, it is important to note that there is nothing inconsistent about accepting both functionalism and the view defined by the latter notion of generalism. Now the reason that generalism *seems* to follow from functionalism is that according to functionalism, mental states are defined in terms of their functional role. Since perceptual states are a type of mental state, and since the functional role of any given mental state type does not depend strictly on which specific stimulus causes it, it follows that perceptual states are more general than the particulars that cause them. However, though this line of reasoning is cogent, it is not contradictory to endorse both functionalist and particularism by defining perceptual states in terms of the specific particulars which cause them. In other words, the functionalist could posit that, though it is possible for a single type of mental state to result from a causal relationship with multiple distinct particulars, the token mental states that result from causal relationships with each particular count (by definition) as distinct perceptual states.

2.2 Lewis' 'Generalism'

Lewis' (1980) argument is akin to Harman's argument in that they both defend a functional (i.e. causal) account of perceptual states. Hence, Lewis writes, "Visual experience depends on the scene before the eyes, and the subject's beliefs about that scene depend in turn partly on his visual experience. The content of the experience is, roughly, the content of the belief it tends to produce" (pg. 239). On this view, experience is characterized in terms of functional relations between stimuli, content, and behavior. The upshot is that neither Harman's nor Lewis' arguments pertain to the question of whether perceptual states are constituted by property instances or non-spatiotemporally located properties. Rather, both theorists defend a functional view about the nature of mental states, where mental states are defined in terms of their functional roles. Here, I consider Lewis' thesis and argue that it does not commit him to any view about whether perceptual states are constituted by abstract properties.

Lewis approximates his thesis as follows:

If the scene before the eyes causes matching visual experience as part of a suitable pattern of counterfactual dependence, then the subject sees; if the scene before the eyes causes matching visual experience without a suitable pattern of counterfactual dependence, then the subject does not see. (pg. 245)

This idea can be elucidated by returning to the teacup example in the following thought experiment: Suppose that at t_1 I am perceptually related to a particular cup₁ that is sitting on my desk. Now suppose also that at t_2 , cup₁ is removed without my knowing, and I hallucinate in its place a phenomenally identical cup₂. My experience has not changed,

and yet I went from being perceptually related to cup₁ at t₁ to not being perceptually related to any cup at t₂. Intuitively, I am seeing at t₁, but not at t₂. Lewis' proposal is this, since my phenomenal experience did not change at t₂, the moment I started hallucinating, it follows that my visual experience did not counterfactually depend on my relation to cup₁ in the sense necessary to say that I 'saw' cup₁ at t₂. On the other hand, leading up to t₁, my visual experience did counterfactually depend on my relation to cup₁ in the sense necessary to say that I 'saw' cup₁. Prior to t₁, I did not hallucinate any cup, and hence, if cup₁ were to be removed, I would cease to experience it. Defining 'seeing' in this way tends to fit with ordinary intuitions about seeing.

This proposal in no way commits Lewis to the view that perceptual states are constituted by non-spatiotemporal properties. Hence, despite the fact that Schellenberg cites Lewis is cited as a canonical generalist, he does not make any claims which fit the definition of generalism most relevant to Schellenberg. (Recall that Schellenberg's only strong argument against generalism depends on the fact that generalism claims perceptual states are wholly constituted by general abstract properties, rather than property instances.) Rather, according to my proposed definitions of particularism and generalism, Lewis' argument commits him to a form of particularism! This follows from Lewis' postulation that, despite the fact that my experience at t₁ is functionally and phenomenally identical to my experience at t₂, I nevertheless count as 'seeing' in the first instance and not 'seeing' in the second. The relevant question here is not whether our perceptual states are constituted by abstract properties, but rather, whether we define perceptual states in terms of the particulars that cause them.

Zooming out, my overall proposal is this: Schellenberg's arguments against generalism (and for particularism) fail to provide any conclusive arguments for particularism, as generalism can explain all of the same phenomena that particularism can. Hence the reason that Schellenberg puts forth the particularity argument, which posits that perceptual states are, by nature, constituted by the particulars they are related to. This is not a reason or a piece of evidence for particularism, but rather an intuition about what the word 'perceive' means. Nevertheless, to bolster her argument, Schellenberg attributes to generalists the view that perceptual states are constituted by abstract properties. But, again, this misses the point entirely, and (as we saw with both Harman and Lewis) none of the people Schellenberg cites as generalists make the claims she attributes to generalism. In the next section, I start from the assumption that what is actually at issue for Schellenberg (and for those engaged in a dispute between so called 'particularism' and 'generalism') is whether the mind-independent particulars which are causally responsible for our perceptual states are constitutive of those perceptual states. I conclude that the actual debate between particularism and generalism is a metalinguistic negotiation. The

upshot is not wholly negative; metalinguistic negotiations need not be trivial, as the way we use words can have significant effects. There is thus good reason to take a pragmatic stance on the nature of perception.

Part Three

Here I argue that the debate, as characterized above, is a metalinguistic dispute. This, in turn, provides a reason to think the dispute between particularism and generalism might be of some practical consequence. But what are ‘negotiations’ and what makes them ‘metalinguistic’?

3.1 Defining Metalinguistic Negotiations

According to Plunkett (2015)—who introduces the concept of ‘metalinguistic negotiation’ to argue that much of philosophy involves normative disputes about what words should mean—a ‘dispute’ is “a linguistic exchange that appears to express a disagreement” (pg. 835). Importantly, not every ‘dispute’ is ‘disagreement’, as disagreement “[...] involves a kind of rational conflict of mental states” whose paradigm case is “of two people holding rationally conflicting beliefs, such as, most straightforwardly, when one person believes a proposition that the other person believes the negation of” (pg. 835). The crucial point is that though disputes appear to express disagreement (i.e. a rational conflict in mental states), they need not actually express disagreement. This accounts for cases where interlocutors ‘talk past’ one another; they mistakenly believe that they disagree.

Now, according to the latter definition, disagreement can take many forms and can be expressed with many different types of linguistic exchange, canonical and metalinguistic disputes being the two major categories. The term ‘canonical dispute’ refers to “any dispute that centers on the content literally expressed by the speakers,” whereas the term ‘non-canonical dispute’ refers to any dispute which “center[s] on information that is communicated via pragmatics, rather than semantics” (pg. 836). This distinction relies on the blurry boundary between semantics, the field of study thought to account for the literal content of ‘what is said’, and pragmatics, the field which accounts for contents which are implied or implicated, but not said.

Because the boundary between semantics and pragmatics is fuzzy and controversial, so too is the distinction between canonical and non-canonical disputes. Nevertheless, there are some clear paradigm cases. Consider (I) and (II).

- (I) A says, “C is tall, he is almost 7 feet!”, and B responds, “C is not tall, he is only 6 feet.”

- (II) A says, “Wow, C is tall, he can reach the highest book on the shelf!” and B responds, “No, C is not tall, everyone but you can reach the highest book on the shelf.”

In both (I) and (II), A and B are engaged in a dispute, as they seem to express disagreement. Further, in both cases, there seems to be genuine disagreement, as either the literal or non-literal content of A and B’s utterances evidences a rational conflict in their mental states. Where (I) and (II) differ is the content they disagree about, and the way that the disagreement is expressed. In (I), A and B disagree over whether C meets the contextually salient threshold for being tall. A believes that C is almost 7 feet, whereas B believes that C is only 6 feet. Their exchange is a canonical dispute, as the disagreement is clearly expressed in the literal semantic content A and B’s utterances. By contrast, in (II), A and B disagree not about C’s actual height (they both know he can reach the highest book on the shelf), but rather, about where they should set the contextual standard for tallness. This dispute is non-canonical, as the disagreement is not expressed by the literal semantic content of A and B’s utterances, but by what is implicated or implied. Further, (II) is a specific type of non-canonical dispute called a ‘metalinguistic dispute’, where “a speaker uses a term (rather than mentions it) to express a view about the meaning of that term, or, relatedly, how to correctly use that term” (pg. 834). In (II), A and B use the term ‘tall’ (rather than mention it) to express conflicting views about how the contextually variable standard for ‘tallness’ should be set.

Before arguing that the dispute between particularism and generalism is a metalinguistic negotiation, it’s important to note one last distinction—the distinction between a ‘negotiation’ and an ordinary ‘dispute’. Negotiations are disputes about a normative (rather than factual) disagreement. Using (II) as an example, it is plausibly the case that A and B are arguing about how the word ‘tall’ *should* be used. However, it could also be argued that A and B are arguing about how ‘tall’ is *in fact* used. The question here is whether the participants in the dispute disagree about what *is* the case, or rather, about what *should* be the case.

3.2 The Metalinguistic Dispute Between Particularism and Generalism

Above, I draw from Plunkett (2015) to make the following crucial distinctions:

1. Disputes vs. disagreements.
2. Canonical vs. non-canonical disputes.
3. Metalinguistic non-canonical disputes vs. other non-canonical disputes.
4. Non-canonical disputes vs. non-canonical negotiations.

Accordingly, I use each of these distinctions to argue that the dispute between particularism and generalism—characterized not in Schellenberg’s terms, but according to my more charitable reading—is a non-canonical metalinguistic dispute that expresses a genuine disagreement about how the term ‘perceptual state’ should be used. In other words, the dispute is a metalinguistic negotiation.

For the sake of clarity, I here reiterate my characterization of the debate between particularism and generalism by returning once more to the teacup example. Consider the following two cases. In the first case, I am perceptually related to cup_1 . In the second case, I am perceptually related to a phenomenally identical cup_2 . Call my perceptual state in the first case ‘ state_1 ’, and in the second, ‘ state_2 ’. In both cases, my phenomenal experience is identical. I take it that the debate between particularism and generalism is a debate about whether $\text{perceptual state}_1$ and $\text{perceptual state}_2$ are distinct states, or identical states. This amounts to the question of whether $\text{perceptual state}_1$ and $\text{perceptual state}_2$ are identical since (though they are caused by distinct particulars) they have the same phenomenal and functional properties, or distinct, because they are partially constituted by the distinct particulars they are causally related to.

Now imagine the following exchange, represented in (III), between a generalist (X) and a particularist (Y).

(III) X: “ State_1 and state_2 are identical perceptual states, because though they are caused by distinct particulars (cup_1 and cup_2), they nevertheless have all of the same phenomenal and functional properties.

Y: “No, you are wrong. Because state_1 and state_2 are caused by distinct particulars, they are distinct perceptual states, regardless of whether they have all of the same phenomenal and functional properties.”

The first question, corresponding to the first distinction between disputes and disagreements, is the following. Does (III) *actually* or merely *seem* to express a disagreement between X and Y? In other words, does (III) evidence a rational conflict in the mental states of X and Y? The answer is yes. Though X and Y seem to agree on all of what Plunkett calls the ‘object-level issues’, there is good evidence that they disagree on an important ‘representational-level issue.’ Roughly speaking, object-level issues are questions about reality itself, as opposed to representational-level issues, which are about representations of reality. Short of providing a strict definition, Plunkett gestures towards the following types of issues as being paradigmatic ‘object-level issues’ of philosophical discourse:

(a) 'ground' refers to a constitutive (rather than causal) explanatory relation between facts, (b) 'supervenience' can be defined in terms of the following: the A facts supervene on the B facts when there cannot be a difference in the A facts without a difference in the B facts, and (c) 'real definition' concerns a definition of what something really is (which, on some views, might centrally involve giving an account of its essence). (pg. 829)

In (III), X and Y may take themselves to be involved in an argument about an object level issue about 'ground' or 'real definition', according to the latter definition. However, since X and Y agree on all of the facts about reality—they agree that state₁ and state₂ are phenomenally and functionally identical, and that state₁ and state₂ are caused by distinct cups (cup₁ and cup₂, respectively)—this interpretation must be wrong. Nevertheless, they do express a genuine disagreement about a representational level issue. Namely, they plausibly disagree over the meaning that 'perceptual state' either *does* or *should* have. This is evidenced by the fact that they each consistently use the term to mean something different from each other. So, again, the answer to the first question is yes; X and Y are not merely talking past each other, and their dispute expresses an actual disagreement. But their disagreement is not about an object level issue (as they might take it to be), but a representational level issue.

The second question, corresponding to the distinction between canonical and non-canonical disputes, is whether the disagreement in (III) is expressed via the literal semantic content of X and Y's utterances, or if it is implicated or implied via pragmatic processes. Admittedly, the semantics/pragmatics distinction is notoriously fuzzy, and so the question of whether the disagreement in (III) is canonical is open to debate. However, since most hearers presumably take the utterances in X and Y to be literally about object-level issues, there is good reason to believe that the dispute about representational issues is facilitated by pragmatic processes. A further reason to believe that the disagreement is non-canonical is that the disagreement has to do with the meaning of 'perceptual state,' but X and Y use the word, rather than mention or quote it. Thus, the disagreement about word meaning is implied, and not explicitly stated. Taken together, these two considerations yield strong support for the conclusion that the dispute in question is non-canonical. These considerations also yield an affirmative answer to the third question of whether the dispute is metalinguistic. It is a metalinguistic dispute because the interlocutors use, rather than mention the word 'perceptual state' to communicate views about the meaning that the term should/does have.

The final question—corresponding to the fourth distinction, between disputes and negotiations—is whether (III) expresses a dispute or a negotiation. In other words, the

question is whether the issue is about the actual meaning of ‘perceptual state’ or about what ‘perceptual state’ *should* mean. Though matters are a bit fuzzy here, I suggest that the dispute is a negotiation because, though one goal of the debate might be to capture people’s intuitions about perception, I do not believe that Schellenberg, or anyone else engaged in the debate, would care whether the majority accepts their definition. The main goal is to produce an account of perception which explains the roles it plays in our everyday lives, and which is a useful scientific concept.

Conclusory Remarks

I conclude that the debate between generalism and particularism (as I characterize it), is a metalinguistic negotiation over the meaning of the phrase ‘perceptual state’. Further, as I take lengths to argue in part one and two of this paper, my characterization of the debate is an accurate portrayal of what is actually at issue for Schellenberg and anyone else engaged in the debate. Though Schellenberg attributes some uncharitable views to generalists, there is little (if any) evidence that they actually hold the views she attributes to the forms of generalism that are challenged by her particularity argument.

The question now becomes: given that the debate between generalism and particularism is a metalinguistic negotiation, why should anyone care about it? Here, I suggest that something’s being a metalinguistic negotiation does not imply that the debate is trivial. Additionally, even if the metalinguistic negotiation between generalists and particularists is trivial, the conclusion that it is trivial is not itself trivial.

A common reaction to learning that some dispute is a metalinguistic negotiation is to conclude that the dispute is trivial, akin to disputes over whether hot dogs are sandwiches or whether tomatoes are fruit. But this conclusion does not follow. Many metalinguistic negotiations have drastic and hugely important consequences. Consider, as plausible examples of metalinguistic negotiations, debates over the meaning of terms like ‘assault,’ ‘freedom,’ ‘free speech,’ ‘implication,’ ‘logical consequence,’ ‘sex,’ ‘gender,’ ‘race,’ etc. Because we are non-ideal agents, the meanings we associate with specific lexical items have drastic practical consequences. This results from the fact that many terms have cultural resonance and baggage which goes above and beyond their semantic content. Further, use of certain lexical items licenses certain inferences. Hence, definitions of the term, ‘assault’, for instance, have widespread consequences. Broader definitions of ‘assault’ will result in the condemnation of a wider range of acts (and visa versa for narrower definitions). Analogously, since people presumably attribute higher epistemic value to perceptual states than states of hallucination, defining ‘perceptual states’ in such a way that the term only picks out epistemically reliable mental states has some arguably significant consequences.

Nevertheless, if you are unconvinced that metalinguistic negotiation over the term ‘perceptual state’ has any significant consequences and still believe that the debate is trivial, then you have nevertheless still learned something useful; namely, *that* the debates between generalism and particularism—and if Plunkett (2015) is right, most (if not all) metaphysical debates which purport to be about object-level issues—are trivial. The ability to identify triviality is a crucial skill of any analytic philosopher and any serious truth seeker. We should not toil away at pointless debates of no practical consequence.

Citations

Harman, G. (1990). The intrinsic quality of experience. *Philosophical Perspectives* 4:31-52.

Lewis, D. (1980). Veridical hallucination and prosthetic vision. *Australasian Journal of Philosophy* 58 (3):239-249.

Plunkett, D. (2015). Which Concepts Should We Use?: Metalinguistic Negotiations and The Methodology of Philosophy. *Inquiry: An Interdisciplinary Journal of Philosophy* 58 (7-8):828-874.

Schellenberg, S. (2018). *The Unity of Perception: Content, Consciousness, Evidence*. Oxford: Oxford University Press.