

The Contents of Consciousness: Reply to Hellie, Peacocke, and Siegel

David J. Chalmers

I would like to thank Benj Hellie, Chris Peacocke, and Susanna Siegel for their very interesting commentaries on *The Character of Consciousness*. All of them focus mainly on issues from the second half of the book, especially issues concerning the contents of consciousness. Hellie focuses especially on the role of acquaintance and perceptual attention in perception and introspection and on the ontology of consciousness. Peacocke focuses especially on my Fregean account of the content of spatial experience. Siegel focuses especially on my account of the Edenic content of perception and its relation to Fregean content.

I will reply to the three commentaries in alphabetical order. Each of the commentaries raises many issues, and space precludes me from replying in detail on each issue. Instead, in each part I concentrate especially on a single central issue: acquaintance with phenomenal states for Hellie, spatial externalism for Peacocke, perceptual Frege cases for Siegel. I aim to discuss the central issue in depth and then discuss other issues more briefly. I have tried to start each section with a relatively accessible discussion of the central issue that does not presuppose much knowledge of my book or of the commentaries.¹

1 Reply to Hellie

Benj Hellie combines a critical discussion of my work with advocacy of a interesting and distinctive positive view. For reasons of space and competence, I address only the critical discussion. I focus mainly on the issue on which he spends the most time: acquaintance with phenomenal states.

¹I am grateful to Benj Hellie, Adam Pautz, and Susanna Siegel for comments on this reply. My reply to Hellie is indebted to an earlier and longer version of his commentary, “Consciousness: From Carnap to Chalmers”, available online at <http://individual.utoronto.ca/benj/ccs.pdf>.

Hellie's discussion of acquaintance concentrates on three main theses that he attributes to me, as follows.

- (A) Subjects are acquainted with the phenomenal properties they instantiate.
- (B) The only property-instances that a subject is acquainted with are their phenomenal property-instances.
- (C) The content of phenomenal properties falsely self-ascribes acquaintance with Edenic properties.

The first two theses are aspects of Hellie's thesis (4a), while the third is his thesis (4f). I accept thesis (A). I do not quite accept theses (B) and (C), but I accept theses that are near enough for Hellie's discussion to be relevant. I discuss all three in turn.

Hellie's central objection to thesis (A) is a phenomenological objection from transparency. There are many transparency theses, but Hellie appears to endorse a particularly strong version: the only properties we can attend to in experiences are properties attributed to external objects. On this view, attention to phenomenal properties is impossible. Given that acquaintance with a property makes possible attention to that property, it follows that we are not acquainted with phenomenal properties.

Hellie asserts his transparency thesis as if it is obvious from phenomenological reflection or at least well-established, but I am not sure of his grounds. The phenomenological tradition largely rejects the thesis for a view on which we have a background awareness of our awareness of objects. G.E. Moore, whose "The Refutation of Idealism" (1903) is the central source of the 20th-century discussion of transparency, denies Hellie's version of the thesis by holding that although our awareness is primarily directly at properties such as blueness, we can become aware of our awareness of blueness "if we look attentively enough". More recently, representationalists such as Gilbert Harman (1990) and Michael Tye (2002) have argued for transparency theses, but their theses are much weaker than Hellie's. Harman argues only that we cannot attend to intrinsic qualities of our experience, but he allows explicitly that we can attend to other features such as intentional features. Tye argues mainly that we cannot attend directly to our experience, in that we always attend to our experience by attending to qualities of the external world.

Some regimentation may be useful here. The general form of a transparency thesis is that it is (i) difficult/impossible to stand in (ii) a certain relation to (iii) certain mental properties or entities. Perhaps the most important dimension of variation involves (iii): this may involve intrinsic

qualities, the awareness relation, phenomenal character, mental states in general, and so on. The relation in (ii) may be attention, direct attention, awareness, introspection, and so on. The dimension in (i) can involve at least difficulty (call this *prima facie* transparency) or impossibility (call this absolute transparency).²

Using this regimentation, we can talk of the *prima facie* transparency of awareness to attention (Moore), the absolute transparency of intrinsic features to attention (Harman), the absolute transparency of phenomenal states to direct attention (Tye), or the absolute transparency of phenomenal states to attention (Hellie). Other transparency theses might include the absolute transparency of experiences to thing-awareness (Dretske), and the absolute transparency of the self to introspection (Hume), and the transparency of belief to direct introspection (Evans). Perhaps one can assume that “absolute” and “attention” serve as default values, so that we can simply talk of Harman’s thesis as the transparency of intrinsic features, Moore’s thesis as the *prima facie* transparency of awareness, Hellie’s as the transparency of phenomenal character (or just phenomenal transparency), and so on.

Regarding Hellie’s phenomenal transparency thesis: Hellie says that he does not find phenomenal attention in his experience, but I think instances are not hard to find. One case involves attention to phenomenal blurring, the shift from clear to blurry vision that takes place when one squints, for example. During this process one clearly is aware of and attends to a change. While one attends to an external object here, one does not represent it as changing. One plausibly attributes fine-grained properties to it at the first stage and coarse-grained properties at the second stage, but these properties are entirely consistent with one another, so no change is represented. Rather, the change is entirely in one’s phenomenal properties: from fine-grained experience (involving awareness of fine-grained properties) to coarse-grained experience (involving awareness of coarse-grained properties). I think that one’s awareness of blurring is clearly awareness of this phenomenal change, and likewise, one’s attention to blurring is clearly attention to a phenomenal change.

Some have used the case of blurriness to argue against representationalism. I do not think that case succeeds: the representationalist can easily account for the phenomenon in terms of representation of fine-grained and coarse-grained properties. But as A.D. Smith (2008) has noted, even if the case against representationalism fails, blurriness yields a compelling case against various transparency theses. These include absolute phenomenal transparency theses. The upshot is that

²Amy Kind (2003) uses “strong” and “weak” transparency for absolute and *prima facie* transparency. Daniel Stoljar (2004) distinguishes many transparency theses, focusing especially on varieties of directness.

even representationalists should deny these theses, holding that one can attend to and be aware of representational properties of one's experience.

A similar moral can be drawn from cases involving shifts in attention.³ In a normal perceptual mode, look around a still room. Nothing seems to change. Now switch to an introspective mode, and look around again. One is now quite clearly aware of a change: for example, when objects disappear from sight. But nothing in the world seems to change. (If one objects that one's head or eyes seem to move, switch the case to a covert change in attention.) The change one is aware of is a change in one's experience: a phenomenal change. I think that likewise, when in the relevant introspective mode, even before looking around the room, one is aware of phenomenal features of one's experience.

So Hellie's phenomenal transparency theses are false. I think the evidence instead tends to support a view on which (i) our perceptual phenomenal states consist in phenomenal awareness of properties (such as Edenic redness) attributed to the external world, (ii) our core acts of attention are to these properties, and (iii) we can also attend to our awareness of these properties, thereby attending to our phenomenal states.

This view is inconsistent with Hellie's transparency thesis, but it is quite consistent with the weaker theses put forward by Moore, Harman, and Tye. In fact, small supplements to this view yield views that entail all three theses. For Moore's thesis, one can add the claim the attention in (iii) is relatively difficult. For compatibility with Tye, one can add the claim that the phenomenal attention in (iii) is always indirect. For compatibility with Harman, one can add the claims that phenomenal awareness of external properties is not itself an intrinsic property of an experience and that we cannot attend to nonphenomenal properties of an experience. I think that claims in the vicinity of these additions are all reasonably plausible. In the case of Tye, the plausible claim is that phenomenal attention always involves externally-directed attention (though I think it is far from obvious that it is mediated by that attention). In the case of Harman, the plausible claim is that we cannot attend to nonrelational (or perhaps nonintentional) features of our experience (though I think it is far from obvious that the relevant intentional and relational properties are extrinsic).

These claims about attention to phenomenal properties do not establish that we are acquainted with those properties, but they suffice to rebut Hellie's main argument against that claim. One way to make the further step to acquaintance is to argue that our capacity for occurrent and transient

³Here I am indebted to conversations with Geoffrey Lee, who has discussed these cases in unpublished work.

attention to phenomenal properties itself requires explanation, and is best explained by a background and constant acquaintance with those properties. I think the most plausible line here is that phenomenal awareness is an acquaintance-involving relation by its very nature: in virtue of the nature of awareness, to be aware of x entails being acquainted with one's awareness of x .⁴

This point bears on another objection Hellie makes to the claim that we are acquainted with phenomenal properties. He worries that acquaintance seems an idle wheel, in that the Lagadonian (reference-constituted) nature of direct phenomenal concepts is enough to do the relevant epistemological work. I reject Hellie's epistemological claim: Lagadonian structure alone explains only why the relevant phenomenal beliefs are always true, not why they are justified. But in any case I first introduce acquaintance for a different purpose: to explain why we are in a position to form Lagadonian concepts of phenomenal properties but not other properties. Hellie suggests that phenomenal attention might also ground these concepts, but something that comes and goes like phenomenal attention will not plausibly be basic here, and our capacity for attention is itself explained by acquaintance.

While I endorse thesis (A) in *TCC*, I do not endorse thesis (B). That is, I say that subjects are acquainted with their phenomenal properties, but I do not say that subjects are acquainted with only their phenomenal properties. Hellie quotes me as saying (on p. 291) that it is a conceptual truth that one is acquainted with exactly one's phenomenal property instances, but in the relevant passage I say only that it is a conceptual truth that to have a phenomenal property is to be acquainted with it. In fact I am committed to denying that the stronger claim is a conceptual truth. Although I think that in the actual world we are acquainted only with Edenic universals, I say in the book (pp. 411-12) that there are Eden-style possible worlds in which subjects are acquainted with Edenic property instances.

Hellie says that his thesis is essential to my view: if someone were acquainted with properties of other entities those entities would seem to be part of them. But that seems wrong: that an Edenic

⁴This is a relative of higher-order representation theories of consciousness, and especially of the Brentano-style self-representational views of consciousness that have become popular in recent years (see e.g. Kriegel and Williford 2006). The key difference with standard versions is that I understand the background awareness as Russellian instance-acquaintance rather than as a standard form of representation (this immediately avoids all objections from higher-order misrepresentation as well as from oversophistication). My version of the view is also nonreductive, in that the awareness relation is irreducibly a phenomenal relation. Of course someone might attempt to turn this into a reductive theory by identifying the awareness relation by a relation understood in functional terms, say. But just as in the case of first-order representationalism (discussed in chapter 8 of *TCC*), this move requires an additional and independent functionalism about the phenomenal, a view that is no more plausible here than elsewhere.

subject is acquainted with a ball's perfect redness does not entail that the ball seems to be part of them. It is natural to hold that acquaintance with phenomenal property instances involve a *de se* mode of presentation, so that it is a nonconceptual analog of *I have property R*, whereas acquaintance with Edenic property instances involves a demonstrative mode of presentation analogous to that in *That has property E*. Perhaps in Eden, Eve could even be acquainted with Adam's phenomenal properties under a demonstrative mode of presentation without those properties seeming to be her own.

Still, it may be that (B) is true in the actual world. It is not obvious that there are counterexamples. I am inclined to think that we are acquainted with numerous nonphenomenal universals, not least including Edenic properties as well as various mathematical and nomic properties, but these cases do not involve acquaintance with instances. Possible nonphenomenal cases of instance acquaintance include temporal properties, on a view where we are acquainted with the time or duration of an experience, resemblance properties between experiences (as Adam Pautz has suggested), and the relation of acquaintance itself, on a view where we are acquainted with acquaintance. But one could argue that in these cases acquaintance with phenomenal properties is what is fundamental. So I am at least open to Hellie's thesis.

Thesis (C) says that the content of phenomenal properties falsely self-ascribes acquaintance with an Edenic external-world property. I do not endorse this thesis in the book. I argue that experiences have Edenic contents that say falsely that an Edenic property is instantiated in one's environment, but I suggest tentatively (pp. 409-10) that these contents do not say that one is acquainted with the property. I speculate (p. 411) that it seems to us introspectively as if we are acquainted with Edenic properties, but I do not say that this seeming is part of the content of experience. Rather, I suggest that phenomenal redness simply involves representing primitive redness (or visually/phenomenally representing it under a rigid mode of presentation; I omit the qualifications for simplicity in what follows).

Still, Hellie's main objection to (C) can be turned into an objection to the view I accept, as follows. In me, phenomenal redness is the property of representing primitive redness. In Eden, phenomenal redness is the property of being acquainted with primitive redness. I know the nature of phenomenal redness, so it does not have a hidden essence. But my Edenic counterpart and I have exactly the same view on things. So the difference in the identity of phenomenal redness is hidden to us. So the nature of phenomenal redness is not revealed after all.

This is unquestionably an interesting argument. I am inclined to respond by denying the second premise. Phenomenal redness is the same property everywhere: the property of representing

primitive redness, or perhaps the property of being aware of primitive redness. In Eden, this property is grounded in and perhaps realized by acquaintance with instances of primitive redness: it is in virtue of being acquainted with these instances that Edenic subjects are aware of primitive redness. Outside Eden, this property is grounded in acquaintance with the universal primitive redness: it is in virtue of being acquainted with this universal that non-Edenic subjects are aware of primitive redness. Both Edenic and non-Edenic subjects are introspectively acquainted with an instance of the property they have in common: being aware of primitive redness. But they are not introspectively acquainted with the grounds (or the realizers) of this property instance.

This line is a sort of cross-world variant of Mark Johnston's account of perception and hallucination in the actual world: perception involves acquaintance with an instance, hallucination involves acquaintance with a universal, and they have acquaintance with a property in common. Hellie acknowledges a response along these lines but replies that it relies on a dubious capacity to make phenomenological distinctions between theory-laden claims about universals and instances. But no phenomenological distinction is needed here. The phenomenology of Edenic and non-Edenic subjects is the same; it is just the grounds of the phenomenology that comes apart. I also note that this way of doing things avoids Hellie's "zombie" argument in which Edenic subjects are acquainted with primitive redness without having phenomenal redness: on this line, the latter is necessitated by the former.

There are a number of residual issues for this view. One concern is that familiar phenomenal properties such as phenomenal redness are now multiply realizable and are not fundamental but derivative, deriving from more basic phenomenal acquaintance relations. If one wanted to, one could preserve the view that phenomenal properties are fundamental in our world by holding that the relevant property is acquaintance with the universal Edenic redness, which is realized fundamentally in our world and derivatively in Eden. In any case I do not see any obvious bad consequences of multiple realizability and derivativeness here.

A second residual issue is the source of the introspective seeming that we are acquainted with instances of Edenic redness. It is hard to see how this seeming could be grounded in acquaintance with phenomenal property instances, given that the latter involve only awareness of the universal Edenic redness and no acquaintance with its instances. An alternative view is that introspection alone tells us only that we are aware of the universal Edenic redness, perception tells us that Edenic redness is instantiated in our environment, so introspection and perception together tell us that we have a veridical perception of Edenic redness. Given this much, it is perhaps a short abductive step to the judgment that we are acquainted with instances of Edenic redness. On this

view the introspective seeming is not wholly grounded in introspective acquaintance but involves a contribution from perception and inference as well. The introspective seeming here is at best a fairly weak one that has nothing like the apparent certainty that attaches to other introspective seemings, so perhaps a hybrid source like this is to be expected.

A third large residual issue concerns the epistemological role of acquaintance with instances and universals, and whether epistemological force derives from phenomenal properties or from their realizers. In the actual world, we are arguably in a position to have justified certainty in introspective beliefs but not in perceptual beliefs. On my view, introspection involves acquaintance with instances (of phenomenal properties), while perception involves acquaintance with universals (Edenic properties). So it is natural to suggest that quite generally, instance-acquaintance but not universal-acquaintance grounds certainty. But now: in Eden, does acquaintance with Edenic property instances ground certainty that they are instantiated?

If we say no, then it seems that either (1) Edenic experience is phenomenally different from a corresponding non-Edenic experience or (2) perceptual phenomenology does not determine whether there are grounds for perceptual certainty. If we say yes, then it seems that (3) Edenic instance-acquaintance carries no extra epistemological powers relative to universal-acquaintance. None of these three views is completely comfortable. I am most inclined to say that given that Edenic subjects have the same perceptual phenomenology as us, Edenic subjects would be open to skeptical arguments just as we are. So I am inclined to say yes and to accept (3).

The residual puzzle then is to explain the epistemological difference between instance-acquaintance with Edenic properties (in Eden) and with phenomenal properties (in actuality). Perhaps one might accept (3a) that there are two sorts of instance-acquaintance, a phenomenologically manifest sort that has the power to ground certainty and a phenomenologically subterranean one that does not. The alternatives are to accept (3b) that epistemological power of introspection derives from something more than instance-acquaintance, (3c) to deny that there are extra epistemological powers to ground certainty in introspection, or (4) to deny that there is true instance-acquaintance in Eden after all, holding that Edenic subjects are acquainted only with Edenic universals, although perhaps this acquaintance is facilitated by some relation to Edenic redness that falls short of instance-acquaintance. I suspect that (3a) and (4) are verbal variants of the same view, and that something in this vicinity is the most attractive view, but I admit that there are significant puzzles here.

I pass quickly over Hellie's remaining objections for reasons of space. Regarding the self: as I note on p. 139 of *TCC*, I am neutral on Hellie's "soul pellet" view (the substance dualist claim that subjects are fundamental). Regarding phenomenal properties: Hellie thinks that the notion

is obscure, but I think we have a strong and intuitive grasp on the idea that there is something that it is like to be some entities but not others, and that subjects can be similar or different in respect of what it is like to be them. Regarding *de se* self-ascription of phenomenal properties: as an old-fashioned Russellian I think one picks oneself out either by direct acquaintance or as the bearer of certain experiences with which one is acquainted. I do not have much to say about the vertiginous question of why that subject is me, which faces materialists and dualists, and *de se* and non-*de se* theorists of content, alike. Regarding other minds: I think that knowledge of others' phenomenal properties comes from an inference to the best explanation of both physical and phenomenal evidence. I reject Hellie's argument that abduction cannot enable ampliative knowledge of others' consciousness: I think that evidence plus abductive reasoning can defeat a hypothesis consistent with that evidence, and I do not agree that all beliefs for which there can be no defeating evidence are indistinguishable from certainty.

2 Reply to Peacocke

Christopher Peacocke focuses especially on my analysis of spatial experience. I treat spatial experience as analogous to color experience, with a narrow phenomenal character and a Fregean phenomenal content that picks out spatial properties in virtue of their being the normal causes of spatial phenomenal properties. Peacocke rejects the analogy between color experience and spatial experience and argues for an externalist view on which both the phenomenal character and the representational content of spatial experience is constituted by spatial relations that obtain in the perceiver's environment.

Peacocke and I agree that phenomenal character determines representational content. I move from here to internalism about some aspects of representational content, while he moves from here to externalism about some aspects of phenomenal character. In particular, he holds that spatial phenomenal character constitutively involves the perception of spatial magnitudes, and that the spatial magnitudes perceived depend on a perceiver's environment, so that spatial phenomenal character also depends on the environment.

A key case here is Brad Thompson's Doubled Earth (2010), which is very much like Earth except that everything is twice the size. Our functional duplicates on Doubled Earth perceive spatial magnitudes twice the size of those on Earth. On Peacocke's view, the corresponding spatial experiences will differ in phenomenal character from our own. Peacocke allows that there is an internal level of sensational phenomenal character, correspondingly roughly to the two-dimensional distri-

bution of colors and the like in the visual field, that will be shared between inhabitants of Earth and Doubled Earth. But crucially, there is also an external level of representational phenomenal character that will be different.

Peacocke's view of Doubled Earth is counterintuitive. In my experience, the great majority of people think that on Doubled Earth, the phenomenal character of spatial experience will be the same as on Earth. Even those sympathetic with phenomenal externalism often hesitate to extend their externalism to cases like this. Presumably the idea is that things seem twice as big to people on Doubled Earth. But it is not at all easy to get a grip on what the difference in phenomenal character will consist in.

What will it be like to be an inhabitant of Doubled Earth? One potential model is the familiar difference in experiences of similar objects occupying the same visual angle but accompanied by different visual cues. In one case we might experience a red square one meter high and ten meters away, in another cases a red square two meters high and twenty meters away. But cases like this will involve a difference in relative spatial representation: for example, in the apparent distance of the square relative to various cues and relative to the size of one's body. So these are not good models for the case of Doubled Earth in which there is no such difference in relative spatial representation.

Perhaps one can contrive a case in which everything in one's current environment seems twice as big. I might be given what I am told is a size-doubling pill, and then experience my body gradually doubling in size while my environment remains constant; objects in my environment might then gradually be replaced by doubled counterparts. Done right (perhaps this could even be tested using virtual reality technology), at the end of this process I might experience myself as 12 feet tall and objects as twice as large as before. But the phenomenal difference here would still involve and arguably would turn on a difference in relative size representation: objects now seem twice as large as remembered objects yesterday. A prediction: if one were to continue in an environment like this with comparisons to the old environment fading away, any phenomenal difference would fade away just in the case of inverting lenses.

Perhaps Peacocke could hold that relative size representation is not crucial here, and that my duplicates on Doubled Earth are phenomenally different from me in much the same way as in these cases even without the difference in representation of relative size. It would be good to know more, however.

Once one countenances Peacocke's idea that the representation of absolute size is constitutive of spatial experience, a certain sort of skeptical hypothesis becomes coherent. Perhaps everything

in my environment is (and always has been) twice as large, or twice as small, as it seems to be. Indeed, perhaps everything is a million times larger or smaller. Speaking for myself, while I find many skeptical hypotheses at least intuitively coherent, I have a harder time with this one. While there is a strong phenomenological case for relative sizes being present in one's spatial phenomenology, I do not think there is the same case for absolute sizes. Correspondingly, I think that the Edenic content of spatial experience involves relative size but not absolute size.

Furthermore, I think that Peacocke's view is subject to my argument from indeterminacy against phenomenal externalism (pp. 354-55). If a subject is unknowingly doubled in size and moved from Earth to Doubled Earth, it is plausible that the externally grounded content of their absolute size representations will gradually shift. For example, a representation corresponding to "one meter" will gradually move from representing one meter to representing two meters, with an intermediate period of "divided reference" during which it is intermediate between the two. If Peacocke's externalism is correct, one's phenomenology of size will also be intermediate during this period. But it is extremely difficult to see how one could have an "divided" phenomenal state that simultaneously and indeterminately involving the phenomenal character of both seeing one size and seeing something twice the size, and it is highly implausible that such a subject would have such a divided phenomenology. Peacocke addresses this argument briefly, saying that on his account, constant sensational content accommodates intuitions in favor of constant phenomenology. But he does not address the argument *against* varying phenomenology, which applies just as much to representational phenomenal content as to sensational phenomenal content.

Peacocke supports his view of spatial experience in part by an analogy with temporal experience. Here the analog of Doubled Earth is Speeded-Up Earth, in which isomorphic subject and an isomorphic environment evolve twice as fast as in our world. He argues that there is no common aspect of narrow content between these subjects, because there is no "zeroth-dimensional" aspect of time to play the role that two-dimensional narrow content plays in the spatial case. Now, my own view of narrow spatial content is unlike Peacocke's model on which two-dimensional sensational content plays this role: on my view, narrow (Edenic) spatial content is three-dimensional through and through. Narrow (Edenic) temporal content will likewise be one-dimensional through and through. Still, the speeded-up case poses an interesting challenge, not least because compared to doubled subjects, there is a stronger intuition that speeded-up subjects will differ phenomenally from us. So the case and the relation to the spatial case deserves discussion.

Although there is some intuition that speeded-up subjects will have speeded-up experience that differs phenomenally from ours, I think that on reflection this is far from obvious. Geoffrey

Lee has argued (in unpublished work) that in fact the experience of a speeded-up subject will be phenomenally the same as that of an ordinary subject. These experiences will extend over a longer period, of course, but Lee argues that time here does not itself serve as a phenomenal property. If Lee is right (and I am at least somewhat sympathetic), then temporal experience is best regarded as constitutively involving the representation of relative (Edenic) times rather than absolute times, just as on my view of spatial experience. This relative temporal content can serve as a sort of narrow phenomenal content. As in the spatial case, we will still perceive and represent absolute times, but these enter the picture at the level of extra-phenomenal Russellian content that is determined by the way that experience is embedded in the world.

Even if Lee is wrong, and the temporal duration of extended experience plays a constitutive role in temporal phenomenology, one can reasonably deny the analogy between the temporal and the spatial cases. In the temporal case, the intuition of phenomenal difference is plausibly driven by the idea that the duration of an experience plays a constitutive view in its phenomenal character. It might do this as an extra-representational element of phenomenal character, or it might do this by playing a constitutive role in temporal representational content. But in the spatial case, although experiences may have a spatial extent (at least on some materialist views), it is not plausible in the same way that the spatial extent of an experience plays a constitutive role in its phenomenal character, either as an extra-representational element or through a constitutive role in spatial representational content. So the considerations at play in the temporal case are not really at play in the spatial case.

If we hold that with Peacocke that the normal and speeded-up cases differ phenomenally, does this undermine a Fregean view of temporal phenomenal content? The question is delicate. One might hold that these cases involve a difference in Edenic temporal phenomenal content, involving the representation of absolute Edenic temporal properties that may not be instantiated in our world (perhaps because they involve passage and the like), but that are matched by ordinary temporal properties in our world. This view is congenial to the Edenic/Fregean picture. Alternatively, one might hold that these cases involve a difference in content involving the temporal properties that are instantiated in our world. This view is still compatible with the Edenic/Fregean picture: we need only say that the relevant Edenic temporal properties are in fact instantiated in our world, so that they serve as their own best match. So either way there is no threat to the two-stage picture. The impact will simply be that absolute Edenic temporal properties are represented in experience and that the duration of an experience may play a role in constituting its Edenic temporal content.

Peacocke connects his discussion about spatial experience to my discussion of the Matrix, say-

ing that if his phenomenal externalism is correct, subjects in a matrix whose spatial experiences are produced by a matrix cannot be perceiving veridically. I am surprised by how he puts this. Given his phenomenal externalism, I would think that the right thing to say is that subjects who have always been in a matrix (which is the relevant class of subjects in my discussion) cannot have true spatial experiences at all. After all, Peacocke presumably holds that they do not stand in the causal and other relations to spatial magnitudes that are required for perceiving and representing those magnitudes and for full spatial experience. So I would expect him to hold that either (i) that these subjects in a matrix have no visual experience at all, (ii) that they have the two-dimensional sensational core of spatial experience but no experience of three-dimensional space, or (iii) that they have a novel sort of visual experience quite different from ours: schmatial experience, perhaps, involving constitutive relations to properties present in the Matrix. Lines (i) and (ii) tend to suggest an easy reply to the skeptic: we can know by phenomenal introspection that we are having fully representational visual experience, so we know that we are not in the Matrix. Peacocke shows no sign of endorsing such a line, so perhaps it is most natural from him to endorse line (iii). But then there is no obvious reason for him to deny that these subjects are perceiving veridically, as I argue they do.

Perhaps Peacocke just means to be saying that his view entails that *if* such subjects have truly spatial experience, as I think they do (and he presumably denies), then they could not perceive veridically. But if his view is correct, the entailed conditional will be irrelevant as its antecedent will be false; and if his view is incorrect, then his reason to accept the conditional disappears.

Peacocke also offers a response to my Matrix argument that is independent of his phenomenal externalism. One of my premises says that Computational Hypothesis, the hypothesis (seriously entertained by many physicists) that physics is itself computational, is not a skeptical hypothesis. And I say that if this hypothesis is correct, then the implementation of the relevant computational structure makes no difference to whether the hypothesis is skeptical. Peacocke responds in effect by denying the latter claim, and holding that the implementation of the computational structure must satisfy certain strong spatial constraints in order for our spatial beliefs to be true under such a hypothesis: if the computation is not laid out the right way in space, our beliefs will be false.

Here I think that Peacocke's attitude runs contrary to the most common attitude within physics. I think that if physicists discovered a level of computational structure that is more fundamental than spacetime physics, and then discovered that this structure was then implemented in some bizarre way at a lower level, this would not lead them to deny that spacetime exists; they would simply say they we have discovered the nature of spacetime. To some extent we have gone through an

analogous process in relativity and even more so in quantum mechanics, where it appears that the most fundamental physics involves a high-dimensional configuration space from which there is no straightforward relation to ordinary three-dimensional space. In both cases, a very small minority of philosophers and physicists have been led to embrace eliminativism about ordinary space, but by far the majority have remained realists. I suspect that the same attitude would be likely to persist given almost any discoveries about how the computational structure underlying physics is implemented.

The attitude among most physicists and philosophers is naturally explained by an underlying spatial functionalism, on which space is identified with whatever plays the role associated with space in physical dynamics and in causing our experience. But this spatial functionalism also leads naturally to a view on which space is present even in a Matrix. Peacocke's attitude is more naturally interpreted as a spatial primitivism, on which we have a primitive grasp on spatial properties independent of the roles that they play, combined with spatial realism, the view that the properties so grasped are instantiated in our world. In *Constructing the World*, I argue at length that this combination of spatial primitivism and spatial realism is difficult to reconcile with contemporary physics.

My own view has elements of both spatial primitivism and spatial functionalism. Spatial primitivism but not spatial realism is true of Edenic spatial properties, while spatial functionalism and spatial realism are true of ordinary spatial properties. If one takes this view, Peacocke is right that there are nonfunctional constraints on the veridicality of spatial content. But these constraints will arise at the level of Edenic content, and will already be plausibly violated in our relativistic and quantum-mechanical world. At the level of ordinary spatial content, these constraints do not apply, so that these contents can be true in a Matrix world. Either way, our spatial representations are no less veridical in the Matrix than in the actual world.

Peacocke makes three other objections to my view of spatial experience. First, he objects that it is unintuitive to deny that we perceive the sizes and spatial properties of objects. In response: I do not deny this. I hold a fairly standard causal view of the perceptual relation (p. 429), on which we perceive ordinary colors (that is, non-Edenic colors such as reflectances and the like) in virtue of their causing relevant experiences. I hold the same view of our perceptual relations to ordinary spatial properties. Perhaps Peacocke means that we *phenomenally perceive* the sizes and spatial properties of objects, where to phenomenally perceive a property is to perceive it (or perhaps to perceptually represent it) in virtue of phenomenal character alone. If so, I do not see the argument for this strong claim. Peacocke suggests that if we did not perceive spatial relations,

we could not know them; but it is not clear why *phenomenal* perception as defined here should be required for knowledge. After all, we certainly perceive objects in the ordinary sense without phenomenally perceiving them, and it is plausible that we know about objects. Perhaps Peacocke thinks that we stand in a much stronger (acquaintance-like?) epistemological relation to ordinary spatial relations than we do to ordinary objects and colors, one that requires something stronger than ordinary perception; but more would need to be said to argue for this claim.

Second, Peacocke objects that spatial content can be common to different sense modalities. Again, I agree with Peacocke. I use this point in the book (pp. 395-7) as part of the case for rejecting the simple Fregean view of spatial content in favor of the more complex two-stage view. On my view, the common content here rests on shared Edenic spatial content between the modalities. In virtue of this shared Edenic content, there is also shared Russellian content involving ordinary spatial relations, in virtue of those properties matching Edenic spatial relations. Peacocke's objection does not yield any obvious objection to this view. Here and elsewhere, I am inclined to think that many of Peacocke's intuitions about the way that spatial content inheres fundamentally in spatial experience are best captured at the level of Edenic spatial content.

Third, Peacocke argues against my internalism by saying that if experiences have only narrow content, they could only make rational beliefs with narrow content, not beliefs with wide content, and that beliefs with wide content are needed to explain action. But on my view, experiences have both narrow content and wide content, and they rationalize beliefs with both narrow and wide content. As before, phenomenal content is not the only sort of content. I do hold that narrow content has a certain primacy here, both in that it determines wide content (with the aid of the environment) and that rational relations among mental states most fundamentally reflect relations among their narrow contents rather than their wide contents. I also reject Peacocke's claim that only wide content can explain action environmentally described: here we must distinguish environment-constituted contents (which are wide by definition) from environment-directed contents (which are what matter in explaining action and which can be narrow). But there is nothing here that requires denying that there are beliefs with wide contents or that they play a role in explaining action. (Similar points apply to Peacocke's discussion of the neural correlates of consciousness and of the principle of structural coherence.)

Concerning two-dimensionalism: I am open to the Peacocke's suggestions that the intensions I discuss are not the truly fundamental Fregean senses, and that they are grounded in some more fundamental sort of sense. I think that the most fundamental senses involve acquaintance rather than reference rules: a reference rule such as "*I* refers to the agent of this thinking" plausibly

turns on some more fundamental acquaintance-involving senses, perhaps for concepts such as *this thinking*. (I also think that facts about acquaintance can explain the rationality of the introspective judgments that Peacocke discusses.) That said, the fact that intensions are not fundamental does not begin to imply that they are redundant. In semantics as elsewhere, all sorts of interesting explanatory work takes place at the non-fundamental level. For example, I think that the intensional framework allows a more straightforward Fregean response to Kripke's anti-Fregean arguments, and a clearer analysis of many issues involving the interaction of epistemology and modality.

Concerning unity: I am inclined to agree with Peacocke that subjects should be given some role in explaining the unity of consciousness. I think that the interaction between subjects and consciousness is likely to be crucial here. Tim Bayne and I suggest that unity is to be explained in virtue of the most basic states of consciousness being states of what it is like (overall) to be a subject at a time. I take it that this approach is congenial with Peacocke's suggestion that subjects be given a key role.

3 Reply to Siegel

Susanna Siegel poses three interesting challenges to my two-stage view of Edenic and Fregean content. First, she casts some doubt on my model of Edenic content by arguing that there can be Frege cases involving Edenic contents. Second, she argues that my account of Fregean content commits me to an implausible psychology of perceptual belief. Third, she suggests that it commits me to an indirect realist epistemology of perceptual belief.

I start with Frege cases. The question of whether there can be Frege cases in perception is an independently interesting one. In the case of thought, we might say in a very loose sense that a Frege case is one in which a subject has two concepts (*Hesperus* and *Phosphorus*, perhaps) for the same object without realizing this. In the case of perception, we might say in an equally loose sense that a Frege case is one in which a subject has two experiences of the same object without realizing this. In the case of object perception, there will of course be countless Frege cases in this sense: seeing the same object from two sides, in a mirror, and so on. For present purposes, we can tighten up the notion to focus on *properties* rather than objects. We experience properties, such as redness or squareness, when things perceptually seem to have those properties. A Frege case for properties is one in which a subject has two experiences of the same property without realizing this.

Can there be Frege cases for properties in perception? One natural route to Frege cases arises

if one holds that we experience external properties such as surface reflectances in virtue of their causal connections to one's experiences. Then it seems that in principle one could experience the same external property twice in different modalities without being in a position to know this. Within my framework, these will be Frege cases involving the ordinary properties presented in Russellian contents, reflecting the fact that they can be picked out under different Fregean contents. But on my view, these ordinary properties are not involved in *phenomenal* contents, those that supervene on phenomenology alone. We might put this by saying that they are not *phenomenally* experienced.

Can there be Frege cases for properties phenomenally experienced in perception? On my view, the properties phenomenally experienced in perception are Edenic properties. I hold that Edenic properties, unlike ordinary properties, are not presented under the separate Fregean modes of presentation that are required for a standard Frege case. So there is a sense in which I must deny that there can be Frege cases for Edenic properties.

We can be more precise. As Siegel suggests, I think that there are Edenic concepts corresponding to Edenic properties. We can form a direct Edenic concept (analogous to the direct phenomenal concepts discussed in chapter 8 of *TCC*) by attending to an Edenic property represented in our experience and forming a concept of it, where the content of this concept constitutively depends on our perceptual awareness of the corresponding Edenic property. (As in chapters 7 and 8 of *TCC*, I am here using 'concept' for mental representations rather than abstract objects.) We can also possess standing Edenic concepts that persist over time and that do not constitutively depend on specific instances of perceptual awareness in this way.

I hold that Edenic concepts are epistemically rigid, in that an Edenic concept picks out the same Edenic property in any epistemically possible scenario. From this it follows that for any two Edenic concepts a and b that pick out the same property, the identity $a=b$ is true in all scenarios. It follows from this that the identity must be knowable a priori, at least on ideal rational reflection (if it were not, there would be an epistemically possible scenario in which it is false).

This commits me to excluding any Frege cases in which the following four conditions are satisfied: (i) the subject phenomenally experiences the same Edenic property in two different experiences, (ii) the subject forms two Edenic concepts a and b of this property, (iii) the subject engages in ideal a priori reasoning concerning the identity $a=b$, and (iv) the subject does not come to know this identity. There may still be quasi-Frege cases in which these conditions are not satisfied: for example because the subject does not form Edenic concepts of the property or because the subject does not reason ideally.

If Siegel's arguments establish that there are Frege cases in which all four conditions are satisfied, this would require significant revisions in my view. It would follow that Edenic concepts cannot be epistemically rigid, and I would have to introduce a further layer of Fregean contents that pick out Edenic properties. Perhaps a picture of the latter sort could be constructed, but it would lose much of the attraction of the Edenic picture I prefer. So much turns on what these arguments establish.

Siegel's central example of a Frege case involves a cyclist who initially experiences a road as flat, and later becomes tired and is not sure whether the road is flat, even under the assumption that her experience is veridical, because she is not sure whether her feeling of increased resistance is wholly due to her changed energy level or partly due to a change in the road. As Siegel describes the case, the cyclist may in fact be experiencing the road as flat, even though she is unsure that she is experiencing the road this way. So although the flatness she is experiencing is the same property as flatness at the first stage (and also the same property as visually experienced flatness), the cyclist will be unable to determine that these properties are the same. All this applies equally if we take flatness to be Edenic flatness, so we have a Frege case involving Edenic contents.

My first reaction to Siegel's case as described is that the cyclist need not be experiencing flatness at the second stage. Perhaps she is perceiving flatness in the sense that her experience is caused by the road's being flat. But is she experiencing the road *as* simply and precisely flat? The cyclist certainly will not say that the road seems flat to her: perhaps she will say that the road seems somewhere in the vicinity of flat, or flattish. If so, it is natural to take her report at face value and say that she is experiencing the road as flattish. The novice cyclist has plausibly not developed the expertise required to perceptually discern precise flatness through bodily resistance. On this interpretation, there is no Frege case here.

This interpretation is supported by Siegel's own observation that in the case of the tired cyclist, "in the large middle range [between extreme levels of tactile resistance at which the road seems to slope up and slope down], there is not plausibly a unique point at which perfect flatness is revealed". I think it is also implausible that perfect flatness will be revealed at multiple points in that range: after all, the variation with increased resistance from representing flatness to representing upward slope presumably involves a continuous monotonic increase in represented slope. Given this much, the plausible conclusion is that for the tired cyclist, there is no point at which perfect flatness is revealed.

Siegel acknowledges the possibility of coarse-grained contents, but she denies that every such case must be treated this way. She says that it is at least possible that a cyclist could experience

precise flatness without being in a position to know that they are experiencing precise flatness, perhaps because they are new to cycling and do not know how to interpret their own experience. Even if this is not the most natural interpretation of the case as initially described, I certainly allow that situations of this sort (in which a property is experienced without the subject being in a position to know this) are coherent possibilities.

My view can acknowledge at least three ways in which something like this could happen. These three ways correspond to violations of conditions (i)-(iii) above. First, the novice cyclist might merely perceive the flatness of the road (in the causal sense) without phenomenally experiencing it, so that it will take further empirical evidence to correlate experiences of this sort with flatness. Second, the cyclist might be non-ideal at attending to their experience and forming direct concepts derived from it, so that although they are perceptually representing precise flatness they form only a concept of flattishness. Third, the cyclist might be non-ideal at reasoning about experience, so that although ideal reflection would connect their precise concept derived from the experience with their standing concept of precise flatness, the cyclist is not in a position to perform this reasoning. None of these three cases will yield a Frege case in the sense required above.

Siegel might insist that there are cases in which first, the novice cyclist is perceptually representing precise flatness; second, the cyclist is an ideal attender who forms a direct concept of the relevant experience that reflects its precise phenomenal character; third, the cyclist is an ideal reasoner who can know anything that could be known by a priori reflection involving these concepts; but in which the cyclist is still not in a position to know that she is experiencing precise flatness, using a standing concept of precise flatness. In response, I see little reason to believe that there can be such cases, and little cost in denying that there can be such cases.

Siegel notes that while the transition from novice to expert sometimes changes an experience (from coarse to precise, for example) it need not. Sometimes an experience stays the same and the subject simply learns to apply a concept to it. Siegel objects that my view will require that whenever a subject learns to apply a concept in this way, the experience will change from indeterminate to determinate. But this is not quite right: I can accommodate unchanging experiences here in each of the three ways described above: either the subject never experiences precise flatness and empirically learns to associate the experience with flatness, or the subject always experiences precise flatness and becomes more ideal at attending to and reasoning about experiences. Again, none of these possibilities will yield Frege cases.

Siegel also suggests that there can be phenomenally identical experiences of “N+” resistance, one experienced in isolation and the other preceded by an experience of lower “N” experience,

so that it is rational for the subject to connect the isolated N+ experience but not the non-isolated N+ experience with flatness. I think that if it is rational to connect the isolated N+ experience with flatness *based wholly on a direct concept derived from the N+ experience and on a priori reasoning*, it will also be rational to connect the non-isolated N+ experience with flatness in this way. Of course the preceding experience may well provide further empirical evidence that should not be ignored in determining whether the road is flat; but this empirical evidence is irrelevant to the issue of a priori connections between the N+ concepts and one's concept of flatness.

Another potential case suggested by Siegel is an intermodal case involving the experience of flatness. Here I think the same sort of analysis applies. If a subject is initially unable to connect a tactile experience to flatness to a spatial experience of flatness and is later able to connect them, then the case will involve either (i) a change in the experience, so that new properties are represented, (ii) new empirical evidence that connects experiences that could not be connected a priori, or (iii) more ideal attention and concept-formation, or (iv) more ideal a priori reasoning. I think that actual cases in which intermodal associations are developed may involve all four of these factors.

Siegel mentions another potential intermodal Frege case in which a subject has a visual experience as of six flashes, a tactile experience as of six taps, but is unable to connect the two. In versions of the case where the subject has a cognitive deficit, it seems clear that diagnoses along the lines of (iii) or (iv) apply. Given ideal attention and reasoning, the subject could simply count to six in each case and connect the two. In cases involving normal subjects and rapidly presented stimuli, it is plausible that (iii) applies: the subjects are not in a position to form a direct concept of "phenomenal sixness". Either way, there is no Frege case here.

Another potential Frege case mentioned by Siegel is the experience of alcohol in different contexts. Here it is arguable that an expert drinker might represent the experience of alcohol in all three cases; but this is a high-level content to which the Edenic account does not apply. If we restrict ourselves to the low-level experience of flavor, it is plausible that the represented flavors are different in all three cases, so there is no Frege case here.

It is worth mentioning a final potential Frege case that Siegel does not discuss but that is discussed briefly by Hawthorne (2007) and at more length by Wishon (2012). This is a case of color indiscriminability in which a subject experiences the same color in different parts of the visual field but is not in a position to know that the colors are the same. The subject might form two direct concepts *a* and *b* (either phenomenal concepts or Edenic concepts) based on these experiences: then plausibly $a=b$ will be true and the subject will not be in a position to know that

it is true. I respond to these cases by invoking versions of (iii) or (iv) above, holding that non-ideal concept formation or non-ideal reasoning is at play here. Either the subject is not in a position to form precise concepts of the relevant properties (instead forming indeterminate concepts that will yield an indeterminate identity), or if they form precise concepts, these will be concepts that could be connected by ideal reasoning that the subject is not in a position to perform. In effect, our perceptual abilities are fine-grained but our reasoning abilities are coarse-grained, and it is the non-ideal coarse-grainedness of these abilities that yields the phenomenon in question.

Siegel's second major objection concerns the psychology of belief in Edenic and Fregean contents. She argues that I am committed to holding that we believe Fregean contents and that this leads me to a dilemma. Either we believe both Edenic and Fregean contents, which is implausible, or we do not believe Edenic contents, in which case I must accept an implausible "psychological defeat" hypothesis on which in forming normal perceptual beliefs, we discredit the face value of our perceptual experience.

I do not accept the psychological defeat hypothesis. I think that typical subjects have beliefs with both Edenic content and Fregean content. Typical subjects know little of the color science that might lead one to reject belief in Edenic contents, so they have little reason to reject such beliefs. So I take the first horn of the dilemma, but the correct path through the dilemma is subtle.

What is the correct psychological story about perceptual judgment in typical subjects? It is implausible that these subjects form two sets of perceptual judgments, one with Edenic contents and one set with Fregean contents. It is more plausible that when these subjects take an experience (of redness, say) at face value, they form a single perceptual judgment, *That object is red*. We can bring out the truth-conditions of this judgment by considering its inferential role. If the subject were to learn that she was in an Edenic world with an Edenically red object present, she would accept the judgment. What if she learns that there no Edenic properties but the object she is looking at has a physical property that causes experiences of redness? I argue in *TCC* that in such a case, a reasonable subject will have an equivocal attitude, taking her experience to be imperfectly veridical but not perfectly veridical. I think such a subject will take the same attitude to their perceptual judgment about redness, taking it to be imperfectly but not perfectly veridical (or true). So I think that like perceptual experiences, perceptual judgments have both Edenic contents (corresponding to their perfect veridicality conditions) and Fregean contents (corresponding to their imperfect veridicality conditions).

Different stories could be told here. If one does not like associating judgments with both perfect and imperfect veridicality conditions, one could regard the content of the judgment as in-

determinate with the Edenic content serving as one precisification and the Fregean content serving as another. Alternatively, one might take one of these standards of veridicality to be the true standard for the judgment—perhaps imperfect veridicality, since this seems to best track our ordinary standards. Or perhaps different subjects form perceptual judgments with different contents. Importantly, on none of these accounts does the subject reject Edenic contents. On the account on which the subject's judgment has only Fregean content and not Edenic content, we strictly speaking take the second horn of Siegel's dilemma (the subject does not believe Edenic content), but since the subject does not reject Edenic contents, psychological defeat does not follow.

On none of these views is there psychological defeat. In addition, none of them involve separate judgments that an apple is perfectly red and imperfectly red. Some unusual subjects might go on to form such judgments: after all, there is no inconsistency in judging separately that an apple is Edenically red and that it has a property that matches Edenic redness. But that is not the usual case.

Siegel asks about the relevance of the Tree of Illusion and the Tree of Science. The trees are supposed to reflect our (philosophers' and scientists') reasons for believing that we are not in Eden: first, our knowledge that there are perceptual illusions, and second, our knowledge of color science. They are not supposed to yield a psychological or epistemological story about the roots of Fregean content in ordinary subjects. After all, many ordinary subjects have not eaten from the Tree of Science and some may not even have eaten from the Tree of Illusion. It is also worth noting that while eating from the Tree of Illusion is enough to tell one that one is not in Eden as originally conceived, it is not enough to justify one in rejecting Edenic contents in general. The Tree of Science is required for that.

What about those sophisticated subjects who have eaten from the Tree of Science? Here, some may continue to accept Edenic contents even on reflection (as some philosophers do). Some may continue to make perceptual judgments with an element of Edenic content while rejecting those contents on reflection; this view will involve an element of compartmentalization, on which reflective judgments are inconsistent with automatic perceptual judgments. Still others may be sufficiently disciplined that they no longer make any perceptual judgments with Edenic content at all. I think that all three of these cases are possible.

Siegel's Psychological Defeat Hypothesis applies only to the third case on which knowledgeable and disciplined subjects refrain from making Edenic judgments. Siegel argues that refraining from taking experience at face value is difficult, but I think it is certainly possible. In an environment containing known illusions (a hall of mirrors or a virtual reality, perhaps), a knowledgeable

and disciplined subject can plausibly refrain from forming perceptual judgments that take their experience at face value, and can form modified judgments instead. If so, it should be possible to refrain from making Edenic judgments too. In any case, this psychological issue goes well beyond the basic Edenic/Fregean framework. The framework itself makes no predictions about whether this sort of restraint is possible and whether it is common.

All this also bears on Siegel's discussion of principle (*), which says that if S knows that a color experience is veridical by ordinary standards, it is rational for S to believe that the properties attributed by the presentational phenomenology of the experience are instantiated. I accept the simpler thesis that if S knows that an experience is veridical, it is rational for S to believe that the properties instantiated by the experience are instantiated, as long as "veridical" and "properties attributed" are disambiguated in parallel (perfect veridicality and Edenic content, imperfect veridicality and Russellian content). Siegel intends her more technical formulation to force a non-parallel disambiguation here (imperfect veridicality and Edenic content). I am not sure that this forcing succeeds, as I think both Edenic and Russellian properties can be said to be attributed by presentational phenomenology. Under Siegel's intended interpretation, I agree that there will be counterexamples to the thesis (most obviously, counterexamples involving sophisticated post-fall subjects), but I think the thesis is now so technical that there is little cost to denying it.

Siegel's final main point concerns what sort of account of perceptual justification fits my account of perceptual content. Siegel says that the account may suggest an indirect realist account of perceptual justification; she does not take this to be an objection to the account, but she observes that others might. I do not discuss perceptual justification at any length in the book, but in principle my view of perceptual content could be combined with many views about justification.

The Edenic view is certainly compatible with a dogmatist view on which perceptual experience with Edenic content provides *prima facie* justification for perceptual beliefs with the same content. For subjects who eat from the trees, this justification may come to be empirically defeated. Siegel suggests that that tree-eating beliefs supported by Edenic contents cannot defeat Edenic contents, but that does not quite seem right: we could certainly learn that the Edenic picture has internal inconsistencies this way. In any case, experiences will also have Fregean contents, which can also support the tree-eating beliefs. It would also be perfectly coherent to embrace dogmatism about these Fregean contents.

I am not especially sympathetic to dogmatism myself. Dogmatism may capture certain natural views in "folk epistemology" about when we treat beliefs as justified. But I do not think that philosophical reflection yields much support for these views. It is hard to see how a perceptual

experience that is not itself justified has the power to confer justification on a perceptual belief, any more than an unjustified perceptual belief has such power. So I think that when it comes to perceptual justification proper, something more than the dogmatist story is needed.

I do not have a well-worked out view of perceptual justification, but I am at least attracted to an account along the following lines (a version of which is sketched briefly at the end of Excursus 15 in *Constructing the World*). The Fregean content of an experience holds (very roughly) that its cause is its normal cause. So we know that if conditions are normal, the Fregean content is correct. But it is arguable that if we know that normally p , we have at least some justification for believing p (perhaps because we have some default justification for believing that circumstances are normal). So we have at least some justification for accepting the Fregean content of our experiences.

It is true that this story about perceptual justification would also be available to an indirect realist. A version of it would even be available to those who deny that experiences have content, in that it will justify the belief that the normal cause of the experience is present (although for this to yield justification for perceptual beliefs, something like my Fregean view of perceptual belief will be required in addition). In effect, the justification itself requires only a priori reasoning and introspective knowledge. In that sense, although I am not an indirect realist, my account of the justification for perceptual belief has the character of indirect realist accounts.

I note, though, that all this flows not from my account of perceptual content but from largely independent epistemological commitments. In particular, it flows from my tentatively held view that introspective justification is basic whereas perceptual justification is not. So these epistemological issues do not threaten my account of perceptual content.

References

- Chalmers, D.J. 2010. *The Character of Consciousness*. Oxford University Press.
- Chalmers, D.J. 2012. *Constructing the World*. Oxford University Press.
- Harman, G. 1990. The intrinsic qualities of experience. *Philosophical Perspectives* 4:31-52.
- Hawthorne, J. 2007. Direct reference and dancing qualia. In T. Alter & S. Walter (eds.), *Phenomenal Concepts and Phenomenal Knowledge*. Oxford University Press.
- Hellie, B. 2013. Against egalitarianism. *Analysis*.
- Kind, A. 2003. What's so transparent about transparency? *Philosophical Studies. Philosophical Studies* 115:225-244.
- Kriegel, U. and Williford, K.W. (eds.) 2006. *Self-Representational Approaches to Consciousness*. MIT Press.

- Moore, G.E. 1903. The refutation of idealism. *Mind* 12:433-453.
- Peacocke, C. 2013. Phenomenal content, space, and the subject of consciousness. *Analysis*.
- Siegel, S. 2013. Are there Edenic grounds for perceptual intentionality? *Analysis*.
- Smith, A.D. 2008. Translucent experiences. *Philosophical Studies* 115:225-244.
- Stoljar, D. 2004. The consequences of diaphanousness. In M. Escurdia, Robert J. Stainton & Christopher D. Viger (eds.), *Language, Mind and World: Special Issue of the Canadian Journal of Philosophy*. University of Alberta Press.
- Thompson, B. 2010. The spatial content of experience. *Philosophy and Phenomenological Research* 81:146184.
- Tye, M. 2002. Representationalism and the transparency of experience. *Nos* 36:137-51.
- Wishon, D. 2012. *Russellian Acquaintance and Phenomenal Concepts*. Ph.D. dissertation, Stanford University.