

Reply to critics

Jeff Speaks

April 26, 2017

1	Is experience representational?	2
1.1	Misrepresentation vs. Ramsification	2
1.2	Representation vs. sensitivity to instantiation	3
1.3	The link between perception and belief and other attitudes	3
1.4	Pautz's puzzle	4
2	Inversion without misrepresentation	5
3	Distinctions between the senses & attention	8
3.1	Distinctions between the senses	8
3.2	The puzzling absence of Frege cases	8
3.3	Attention and consciousness	9
4	Objects, kinds, and indiscriminable phenomenal differences	10
4.1	Big undetectable phenomenal differences	11
4.2	Contingently or necessarily indiscriminable phenomenal properties?	12
4.3	Indeterminacy of phenomenal properties	14
5	Worries about availability	15

Thanks to all three of my critics for their careful reading of *The Phenomenal and the Representational*, and for their insightful criticisms. Here I want to focus on their objections to my treatment of the following five topics: (1) the thesis that experience is representational, (2) my denial of the possibility of spectrum inversion without representational change, (3) my account of the senses and attention, (4) my view that objects and kinds can be parts of phenomenal properties, and (5) my ‘availability-based’ arguments for that view.

1 IS EXPERIENCE REPRESENTATIONAL?

1.1 *Misrepresentation vs. Ramsification*

I give two main arguments for the thesis that experience is representational in the book, neither of which is especially original. A condensed version of the first is as follows: in illusions we misrepresent the world; when we misrepresent we stand in some representational relation to a false content; in the case of illusion this relation cannot be judgement, or inclination to judge; so there must be some other special representational relation involved in experience, which I call the sensing relation.

This relies crucially on the premise that in illusion we misrepresent the world; we are, in some sense, getting things wrong. Pautz and O'Callaghan worry: what do we say to someone who just denies this premise?

Pautz suggests a different way of getting content into the picture: using the Ramsey-Lewis method, we define a relation to propositions which stands in some tight relation to phenomenal properties. We can define this relation without relying on any putative commonsense assumptions about misrepresentation, and then go on to theorize about the relation so defined. He suggests that much of what follows in the book could be left unchanged.

I agree with this last point and so, to that extent, regard this suggestion as a friendly one. But there are two reasons to doubt that this way of proceeding is preferable to my own.

First, it would be a mistake, I think, to frame the issue between Pautz and me by saying that he relies on a Ramsey-Lewis definition where I rely on intuition. Virtually any way of introducing a theoretical term, including my way of introducing sensing properties, can be cast as a Ramsey-Lewis definition. It is just that my Ramsey-Lewis definition takes facts about when subjects misrepresent as part of the defining condition, whereas Pautz takes constitution of phenomenal properties as part of the defining condition.

Second, this does not remove the possibility that someone might simply deny that there are any representational properties satisfying the Ramsey-Lewis definition. For, as Pautz notes, just as someone might flatly deny that we misrepresent in the case of illusion, an opponent of a representational view of experience might flatly deny that there is any relation satisfying the kind of Ramsey-Lewis definition that Pautz provides. And in fact it will introduce a new kind of bedrock dispute, for, not just opponents but also at least some proponents of a representational view of experience will also deny that anything satisfies Pautz's defining condition. Consider, for example, a representationalist with a view like the one that Levin defends. She will not think that there is any relation satisfying Pautz's definition (since she will think that the same phenomenal property can in different cases correspond to relations to distinct propositions).

It seems plausible that the best view here is a pluralist one; perhaps there is no all-or-nothing answer to the question of which starting point is best. Rather, it comes down to the question of which philosophers one wants one's argu-

ments to engage. Given an answer to this question, one should then pick a Ramsey-Lewis definition whose defining condition those philosophers will take to be satisfied. I suppose I am more interested in engaging with the views of representationalists like Levin than those who simply deny that there is any misrepresentation going on in stock cases of illusion. In part, this is because I do think that our views about when subjects are misrepresenting are genuinely pre-theoretical. When one presents introductory students to the Müller-Lyer or other illusions, they are quick to describe these as cases in which subjects of the illusions are misrepresenting the environment around them. Doubts about whether there is any genuinely pre-theoretic intuition to this effect thus seem to me a bit overblown.¹

1.2 *Representation vs. sensitivity to instantiation*

O’Callaghan suggests an alternative to the view that experience is representational; namely, that experience involves, not a relation to a content, but a non-representational sensitivity to the instantiation of properties in one’s environment.

Against this, I would pose an unoriginal dilemma: what do we say about cases of illusion? Either we say that the sensitivity is still operative in that case, or we do not.

If it is, one might wonder whether this is a terminological variant on a representationalist view. For then the proponent of this sort of view and I agree on the things to which we stand in the important phenomenal-property-determining relation; it’s just that I take this to be a representational relation, and my opponent takes this to be a non-representational sensitivity relation. Absent further information about the distinction between these two sorts of relations, it seems reasonable to wonder whether this dispute is a merely verbal one about how to deploy the term ‘representational.’

If it is not, then we are left without an explanation of the fact (at least I think that it is a fact) that subjects of illusory experiences genuinely instantiate phenomenal properties.

1.3 *The link between perception and belief and other attitudes*

The second argument for a representational view of experience that I give in the book — like the first one, it is not original — focuses on the transitions between experiences and beliefs. In the normal case, we make swift and effortless transitions between experience and beliefs about our environment; the best explanation of these transitions, I suggest, is that in making these transitions I am simply taking a new attitude — belief — toward the very same content to which I earlier stood in a sensing relation.

O’Callaghan, though, notices a bit of bait and switch in the book. In the beginning, I present the argument just sketched; and then later I adopt the view

¹That is not to deny, of course, that one might give arguments for rejecting this intuition.

that in fact perception and belief are relations to different kinds of contents. What gives?

My view is that the difference in content types does not ruin the explanation of the ease of our transitions between experience and belief. On my view, it is roughly the transition between self-ascribing some property F , and coming to believe that I am F . F and the proposition that I am F are distinct — but it is reasonably easy to see how an ordinarily subject might be able to unproblematically move from the self-ascription of the property to the belief that she has the property.²

Here again I think that I am on firmer footing than the sensitivity-type view which O’Callaghan considers as an alternative. For we can pose the same dilemma as in the previous section. Either we have a notational variant on a representational view, or we get no explanation of how we manage to form judgements in the case of illusory experiences, as of course we often do.

Now, there is a point that O’Callaghan makes here that I agree with: the considerations I give in favor of the representational view don’t compel one to have this view. They are just considerations which, in the absence of good arguments in the other direction, should lead us to the representationalist position. The overall evaluation of the position then depends on whether there are problems with the representational view which should lead us to look for alternatives.³

1.4 Pautz’s puzzle

While himself a proponent of a representational view, Pautz raises a very interesting problem of just this sort. The puzzle is this: there seem to be certain ‘laws’ which govern which contents can be contents of experience. For example, one can’t visually represent a shape without representing it as in some orientation relative to oneself. But no such rules seem to govern other representational states, like belief. So if perception really is representational, then we should expect some explanation of these laws. But what could it be?

I think that this is an excellent puzzle. I don’t have a full solution; some sketchy thoughts follow.

An initial reply is that there are some ‘laws’ which govern other representational states of subjects. For example, my 9 year old daughter is currently able to believe a great number of facts about Harry Potter books; but she can’t currently believe propositions about the incompleteness of arithmetic. The natural explanation of this fact is that she currently lacks the requisite conceptual capacities. So why not explain the corresponding laws about perception in terms of our perceptual capacities?

²Here I am setting to the side my own heterodox views about propositions, which are orthogonal to the present point.

³In this sense, I agree with Pautz when he objects that it ‘doesn’t follow’ from the fact that there are these transitions from experience to contentful states that experiences themselves must be representational (note 1). It was intended to be an inference to the best explanation, and abductive reasoning is typically not valid.

Here Pautz has a ready response: the rules which govern which propositions my daughter can believe are contingent, and hence apt to be explained by contingent facts about her conceptual capacities. But the laws of perception that he singles out are metaphysically necessary, and hence cannot be explained in terms of contingent facts about our perceptual capacities.

In reply, I am tempted to doubt whether his laws of perception really are (metaphysically) necessary. Consider the phenomenal differences between representation of shape in visual and tactile experience. Could there not be some other modality of experience — some other way of instantiating a sensing property — which we do not possess, and in which shapes can be represented as in one's environment but without any particular orientation relative to oneself? I am not sure why this should not be possible. (To be sure, it is pretty hard to see how such a capacity could have been evolutionarily useful — but that's another question.)

I'm not at all sure that there could be such alien modalities. But it would be hard to rule them out on the grounds that we can't imagine what it would be like to have experiences in this sort of alien modality, since it seems quite plausible that there could be ways of sensing which creatures with our perceptual capacities can't imagine.

So suppose that we accept the possibility of an alien modality of the sort just sketched. One might still point out that the laws Pautz cites seem to hold with metaphysical necessity of specifically *visual* representation. Surely there could not be a visual experience of an object as having a shape but no particular orientation; and what, we might ask, explains this?

This is to defend the existence of laws of Pautz's sort, but restricted to a given sense modality. Any defense of such restricted laws owes an account of the nature of the modality to which the restriction is made. But such an account might furnish the materials needed for the explanation of the law; we might, for instance, try to derive the laws which govern the propositions which can be visually represented from the definition of visual experience.

So, in sum, my best guess at how to reply to Pautz's puzzle is: deny that there are such metaphysically necessary laws for sensing properties in general; and hope that we might be able to explain restricted versions of these laws in terms of the nature of the specific modalities of experience to which they are restricted.

2 INVERSION WITHOUT MISREPRESENTATION

Levin is on my side in thinking that experience is representational. But she points out that my view makes spectrum inversion without misrepresentation impossible, and suggests an interesting alternative view which would make it possible. I want to briefly discuss her reasons for wanting to make room for inversion without misrepresentation, before discussing the alternative view she proposes.

Her thought — which many share — is that denying that inversion with-

out misrepresentation is possible is unattractive because it seems “arbitrary to charge one set of perceivers, but not the other, with misrepresenting the colors.”

It seems to me that there are three natural ways to cash out this charge of arbitrariness. Let me briefly sketch these, and say why I don’t find them compelling.

First way: one might mean that there is simply nothing that could ground — no fact about the world which could explain — that one is right and the other is wrong. But, in reply (and following Byrne and Hilbert (2007) among others): suppose that one represents the object as yellow and one represents it as blue. Suppose that the object is in fact yellow.⁴ Why isn’t that enough to ground the fact that one of the individuals is right and the other wrong?

Second way: both of a pair of spectrum-inverts are epistemically blameless in forming their views about the world. To pick one as misrepresenting would be arbitrary in the sense that it would convict someone of error who has been, from an epistemic point of view, impeccable. But, in reply: there are plenty of cases, Cartesian skeptical scenarios being perhaps the most dramatic, in which we should be willing to say that epistemically blameless individuals are getting things wrong. Why should we be hesitant to say the same in the present case?

Third way: as Block (1999) and others have pointed out, cases of ‘spectrum shift’ are not just possible, but actual. Surely, the objection goes, it is arbitrary to say that large groups of statistically normal perceivers are systematically getting the colors of things wrong. But, in reply: they are only getting things a little bit wrong, and the fact that they are getting things a little bit wrong is an entirely predictable consequence of the marvelous determinacy of visual representation. If I were forced to form beliefs about the colors of various elements of my wardrobe with the same level of determinacy, I would also get the colors of the clothes that I own a little bit wrong.

So I am a little skeptical about the idea that we should want, all else equal, to make room for inversion without misrepresentation. But let’s suppose for the sake of argument that we do. How does Levin make room for it?

Her argument begins with the plausible thought that in cases of perceptual constancy we sometimes get changes in (broadly speaking) color phenomenology with no change in the color property represented. In such cases, she suggests, the change in phenomenal properties is due not to any change in the way the perceiver represents the world, but rather to non-representational aspects of the experience. This hybrid view seems attractive, because it both lets experiences genuinely represent aspects of the world, while making room for phenomenal variation without changes in how the world is represented. This will especially be the case, Levin suggests, when subjects of experience have certain cues — familiar from discussions of perceptual constancy — which tip off the perceiver that the phenomenal change is due to something in her rather than to some change in the world. Given this, Levin asks, should we really be so sure that psychedelic phenomenology + constant representation of color properties

⁴Here I set aside relativism about color; see Chapter 23 of *The Phenomenal and the Representational* for discussion of this alternative.

is impossible? Perhaps this is just another case of phenomenal variation with constant representation of the colors of things.

I have four main points to make about this position.

First, I don't agree with Levin's characterization of cases of color constancy. For consider such a case, and grant that the subject represents no change in the color of anything. Does the subject represent *any* change in any properties of anything in her environment? If we are just thinking about how things seem to subjects of experience, I think that perceivers would say: 'It seems to me that there is some change in the world — not just a change in me.' For example, it may be a change in the illumination of some surface.⁵ If some view of this sort is correct, then cases of color constancy don't provide an example of phenomenal change without a change in the subject's representation of her environment.

Second, even if Levin's view of color constancy is right, it is at least not obvious how this makes room for the possibility of the kinds of cases of psychedelic phenomenology + constant representation of color properties which play a role in the argument of the book. For in those cases, we can stipulate that there are no changes in illumination of the relevant sort; adding this stipulation to the description of the relevant scenarios doesn't, I think, affect any of the arguments which follow.

In response to this second point, I think that Levin will reply by saying that color constancy is just one kind of case in which we get phenomenal changes due to changes in non-representational properties of the experience. Other cases arise when the subject has information, or some other sort of cue, which leads him to think that the color of the relevant object is constant despite the phenomenal changes he is undergoing. This suggestion leads to my third and fourth points.

The third point is that this sort of position gives a subject's views about what is out there in the world a surprising degree of power to influence the content of her experience. In a Müller-Lyer case, I know that the lines are the same length; but I think that my experience still represents them as different lengths. Just so, in the color blindness case, I find it implausible that the color blind subject's visual representation of the colors of objects changes depending on his awareness of his condition. Suppose, to take a parallel case, that you are given false information that you are color blind, which you then believe — would the contents of your experience change? I doubt it. But then I don't see why cues should have the power that Levin gives them in other cases.

Fourth, even if it is true that the contents of subjects' experiences are affected to this degree by their beliefs about their environment, it is not obvious that this makes room for the inversion scenarios which Levin is concerned to protect. For consider a garden variety inversion scenario, in which we have a pair of yellow/blue inverts looking at a lemon. For the cases on which Levin leans to provide reason to think that both might be representing the lemon as having the same color, it must be that one of the inverted subjects is in possession of perceptual cues or some other sort of information about her environment which

⁵See, e.g., Hilbert (2005).

affects the content of her experience. But which one? In ordinary inversion scenarios (as in paradigm cases of spectrum shift) both subjects are having what seem to them to be completely ordinary experiences. So it is hard to see what fact about one, but not the other, could introduce the non-representational elements of experience which would explain their phenomenal differences without also introducing a difference in the color they are representing the lemon as having. In this sense, Levin's view is as committed to breaking the apparent parity between the subjects as mine.

3 DISTINCTIONS BETWEEN THE SENSES & ATTENTION

O'Callaghan raises a number of questions about my views about the distinctions between the senses, and the distinction between sensing and attention.

3.1 *Distinctions between the senses*

I suggest that in our account of phenomenal properties there is no need to rely on distinctions between the senses, and that what distinctions between the senses we draw depends on our explanatory purposes.

In reply, O'Callaghan sketches an interesting view of the senses as faculties the exercise of which might help to explain the nature of various perceptual episodes. He suggests that this is a way of avoiding my view that what distinctions we should draw between the senses is a merely pragmatic issue.

I agree with O'Callaghan on substance, and disagree only on a point of bookkeeping. O'Callaghan has identified an interesting explanatory task — the task of explaining the nature of various perceptual episodes — and has introduced distinctions between the senses as part of this explanatory project. Relative to this explanatory project, I agree, distinctions between the senses may well be needed. My view is just that these distinctions are not needed for the explanatory project of giving an account of the phenomenal properties we instantiate, and that the question of how many senses there are outside the scope of any explanatory project is not an especially clear one.

3.2 *The puzzling absence of Frege cases*

O'Callaghan's second point on this topic worries me more: he points out that if sensing and attending are really distinct relations to contents, we should expect there to be 'Frege cases' in which we are unsure whether the object to which we are attending is the same as the one we are sensing. But there aren't.⁶ What

⁶To be sure, this is possible when we are sensing an object in more than one sense modality; so I might visually represent an object as at a location and represent a sound as at a location, attend to the location as visually represented, and be unsure whether that is the same location as the one from which the sound is represented as originating. What seems impossible is that a subject sense some property, attend to that property, and be unsure, for *all* of her sensory relations to that property, whether the property is the same as the one to which she is attending.

might explain this?

While I do have idiosyncratic views about attention, it is not clear to me that the question O’Callaghan astutely identifies arises from theoretical commitments specific to my views. After all, almost everyone thinks that we can visually sense something without attending to it, and so recognizes the distinction between sensing and attending to some content. Given this fact, everyone owes some answer to the question of why Frege cases cannot arise when we bear these distinct attitudes to a single content.

However, even if the question arises independently of my views on attention, it may still be that my views make the question harder to answer. So, for instance, O’Callaghan suggests the possibility that attending is an adverbial modification of sensing; attending to something is simply sensing that thing in a certain way. On one way of understanding the underlying metaphysics, this would make the basic representational state undergone by subjects in experience a matter of a three-place relation obtaining between the subject of the experience, the sensed content, and the objects of the subject’s attention. In this sense, it might be understood as at least formally analogous to familiar views of belief as a involving a three place relation between a subject, a proposition, and a mode of presentation. This would be a departure from the kind of view that I defend in the book, though I suspect that most of what I have to say could be re-cast in these terms.

Here is an alternative suggestion, which draws on some recent work in the philosophy of language.⁷ We can distinguish between a series of thoughts about some object *o* where the subject recognizes *o* as the same throughout, and one where she does not. It is plausible that this cannot be explained in terms of commonality of content between the thoughts — that is one moral of well-known examples like the ‘London’/‘Londres’ example of Kripke (1979) and the ‘catsup’/‘ketchup’ example of Salmon (1990). So recognition of an object or property as recurring in one’s representations is a distinct type of psychological state. One might take the moral of O’Callaghan’s observation to be that the right theory of attention will have to build this sort of recognition of recurrence into the account of the nature of the state. How this would go, I admit that I don’t know. O’Callaghan’s point here strikes me as important, and worthy of more thought.

3.3 *Attention and consciousness*

O’Callaghan also poses a dilemma for me about the relationship between attention and consciousness. The dilemma, put in my preferred terms, is this: are the phenomenal properties instantiated by a subject due only to what she attends to, or can sensed but un-attended properties also affect phenomenal properties?

I go for the second horn of the dilemma: subjects can differ in their phenomenal properties even if there is no attentional difference between them.

⁷See especially Fine (2007); Salmon (2012); Soames (2015).

Against this, O’Callaghan points out that if attention is necessary for demonstrative thought, we cannot give any ‘availability-based’ argument for the existence of sensed contents outside the scope of attention, and that we therefore need some other reason to believe in them. If this view of demonstrative thought is correct — something on which I am agnostic — then I agree. However, I think that a reason for positing sensed but un-attended contents is relatively easy to come by. There could be a pair of experiences which differ only with respect to the properties represented outside the periphery of the perceiver’s attention, and the subjects of those experiences could yet differ in their phenomenal properties. If one is convinced by prior arguments in the book that a difference in phenomenal properties entails a difference in contents sensed — though of course not all will be — that gives us the reason we need.

4 OBJECTS, KINDS, AND INDISCRIMINABLE PHENOMENAL DIFFERENCES

One of the stranger aspects of my view — Pautz more gently says ‘unique’ — is that I think that objects and natural kinds can be parts of the contents to which I stand in sensing relations, and hence, on my view, also constituents of phenomenal properties.⁸ This means that two subjects can have different phenomenal properties in virtue of standing in sensing relations to distinct but indiscriminable objects or natural kinds. So these will be genuinely distinct but indiscriminable phenomenal properties.

In the book I argue for this in basically three steps.

First, I give availability-based arguments in favor of the thesis that sensing properties can include relations to objects and natural kinds. Given a general preference for a view which explains the supervenience of the phenomenal on the representational in terms of their identity, this points toward a view on which objects and kinds can also be parts of phenomenal properties.

Second, I suggest that the main reasons for finding my views odd is a general aversion to distinct but indiscriminable phenomenal properties. But, I argue, the phenomenal sorites and some other cases show that there can be distinct but indiscriminable phenomenal properties.⁹ So, since we have an argument for letting in objects and kinds, and no good reason not to, we should think that phenomenal properties involve relations to objects and natural kinds.

Third, I argue that sometimes a claim about some property can seem less counterintuitive once we are given an account of what that property is. An admittedly imperfect analogy is given by the identity of heat and mean molecular

⁸Though I don’t want to exaggerate my originality here; as noted in the book, my views on these topics are very much indebted to Johnston (2004). I should also add that, while it is approximately right to say that I take objects and kinds to be constituents of phenomenal properties, it is more accurate to say that phenomenal properties have contents as constituents which in turn have objects and kinds as constituents. I do not assume that the relevant ‘is a constituent relation of’ is transitive.

⁹I discuss other cases, which involve what I call ‘phenomenal match’ and the representation of change, in Chapters 31 and 32.

motion. Prior to learning of this theoretical identity, the view that there can be distinct but in principle indiscriminable levels of heat might seem strange; given this theoretical identity, it seems quite natural. Much the same, I think, is true in the present case. If one accepts the view that phenomenal properties are certain sorts of representational properties, then, once we see what sorts of things these representational properties are, this makes it easier to see how certain pairs of phenomenal properties could be distinct but indiscriminable.

Pautz questions how good the availability-based arguments mentioned above are. I'm a little worried about these arguments myself, on roughly parallel grounds; I return to these in §5 below. But he also provides interesting independent arguments against letting objects and kinds into phenomenal properties, and the arguments he gives don't rely on the assumption that distinct phenomenal properties must always be discriminable. As I understand him, he gives three main arguments against the view.¹⁰

4.1 *Big undetectable phenomenal differences*

Pautz himself is happy to agree that there can be distinct but indiscriminable phenomenal properties and, like me, takes the phenomenal sorites to illustrate this fact. However, these are cases in which only quite small phenomenal differences are indiscriminable; and, as he points out, my view admits the possibility of quite large but still indiscriminable phenomenal differences. One might find the second sort of possibility objectionable even if one does not have this view of the first sort of case.

The simplest case to illustrate his point is a pair of experiences which differ only with respect to representation of distinct natural kinds K and K^* . Suppose that K and K^* are not just different, but *very* different. Then the representational properties involved in the two experiences will also be very different. And, if phenomenal properties just are representational properties, it follows that the phenomenal properties instantiated by the subjects of the two experiences will also be very different. But they will still be indiscriminable; hence large but indiscriminable phenomenal differences. A parallel case could be run using discriminable but nonetheless very different objects.

An initially tempting but bad reply is to say that the phenomenal properties are not *that* different, since they differ only with respect to representation of a kind property. But then (to pick up on another of Pautz's points) consider a pair of experiences which differ only in that one represents an object as red and the other represents that object as green. The phenomenal properties here would seem to be just as similar as in the case of the kind-experiences, and yet (here I agree with Pautz) the red/green experiences could not involve indiscriminable

¹⁰A worry which Pautz does not press, but which I think is a serious worry, is that the dispute between the proponent of my sort of view and someone who restricts sensed contents to purely general contents is a verbal dispute about how to use phrases like 'sensed contents' or 'content of experience.' I think that this is a possibility to which not everyone working on these topics — I include myself here — has been sufficiently alive. I hope to explore this more in future work.

phenomenal properties.

A better reply involves a distinction that I make in Chapter 29 of the book between two different classes of contents which can figure in sensing properties: the phenomenally silent and the phenomenally loud. Roughly, a type of content is phenomenally silent iff two phenomenal properties which differ only with respect to contents of that sort are always indiscriminable, and phenomenally loud otherwise.¹¹ Given the existence of phenomenally silent contents, any principle to the effect that ‘big phenomenal differences must be detectable’ has to be restricted to differences in the representation of phenomenally loud properties. (This is something about which I should have been much more careful in the early stages of the book, before I introduce the notion of a phenomenally silent content.)

Of course, one might doubt that there *are* phenomenally silent constituents of phenomenal properties; more on this below. But once they are on board, the existence of big differences in phenomenally silent constituents of phenomenal properties without a discriminable difference does not seem to me to be a substantial extra cost of the view.

4.2 Contingently or necessarily indiscriminable phenomenal properties?

A second sort of worry Pautz raises can be presented as a kind of dilemma. Take some putative pair of distinct but indiscriminable representational properties *F* and *G*. First horn: it is metaphysically impossible for someone who instantiates *F* and *G* to be able to discriminate between them. Second horn: they are just contingently indiscriminable.

Pautz suggests that I take the second horn. But I would be very unhappy if forced into this view. For consider experiences which differ only in that they represent distinct but qualitatively identical objects. It is just very hard for me to see how any possible subjects of experiences which differ only in this way could discriminate the two phenomenal properties.

My view is instead that phenomenal properties which differ only in their phenomenally silent constituents are necessarily indiscriminable. (If I had it to do over, I would add a modal operator to the definition of ‘phenomenally silent’ on p. 215, which would make this view more explicit.¹²) Pautz wonders: ‘What could explain distinct phenomenal properties being not just contingently, but necessarily, indiscriminable?’

¹¹Pautz notes, fairly, that the term ‘phenomenally silent’ is not ideal. One might think that it means ‘not making any difference to phenomenal properties.’ But that is not what I mean by it, since I think that these contents can figure in sensing properties, and take phenomenal properties to be identical to sensing properties. It is hard to think of a better non-cumbersome term; perhaps ‘indiscriminable contents’?

¹²The one place in the book where this view does come out is the definition of the relevant sort of indiscriminability (209). As I use it in the book, indiscriminability is a modal notion; for phenomenal properties to be indiscriminable is for it to be impossible for any subjects to know that they are distinct on the basis of introspection and memory.

If this is a pressing question it is, I think, a question which anyone who believes in indiscriminable phenomenal properties — and not just someone with my views about objects and kinds — must answer. For consider a pair of consecutive experiences in a phenomenal sorites in which the subject is intently focused on the experience. It seems to me quite plausible that the phenomenal properties instantiated by the subject at those times are not just contingently, but necessarily indiscriminable. It is very hard for me to see how someone could have just those experiences, and — by virtue of superior introspective capacities — discriminate them.

One might be tempted to reply that this is in fact easy to imagine, since we can just imagine someone with perceptual capacities more discriminating than the subject in the sorites series. But this reply would be confused. A subject with more discriminating perceptual capacities would be having *different* experiences, with different contents, and would be instantiating different phenomenal properties. That is after all the usual effect of improvements in one's perceptual capacities. (Compare: putting on glasses.) The fact that such a subject could discriminate between *her* phenomenal properties hardly implies that some subject with the same phenomenal properties as our original subject could discriminate them.

But this does not answer Pautz's question: why should some distinct phenomenal properties be necessarily indiscriminable? I don't have a settled view on this topic, but one strategy would appeal to the transparency of experience. It is by now a familiar thought that when we try to introspect our experiences, we end up focusing on the objects and properties that our experience presents as in our environment. This is widely taken to be a necessary rather than merely contingent feature of experience. Setting aside non-veridical experiences for simplicity, let's say that a subject can *perceptually discriminate* between two scenes iff that subject can, solely on the basis of her experience, tell which scene is before her. Then the transparency of experience makes the following thesis plausible:

□ (*S* can discriminate between phenomenal properties *F*, *G* iff *S* can perceptually discriminate between the scene before her when she is instantiating *F* and the scene before her when she is instantiating *G*)

No doubt this would need some refining.¹³ But something like it might provide the explanation Pautz seeks. Why is it that sometimes distinct phenomenal properties are necessarily indiscriminable? Because the scenes presented when the subject instantiates those phenomenal properties are perceptually indiscriminable, and (given the transparency of experience) it is impossible to have discriminable phenomenal properties without perceptually discriminable scenes.

¹³One needed refinement would be a stipulation that the subject's perceptual capacities and orientation with respect to the scene be held fixed; another would involve the kinds of attention shifts discussed in Chapter 27.

4.3 *Indeterminacy of phenomenal properties*

Pautz also worries that, on my view, it will sometimes be indeterminate which phenomenal property I am instantiating. Suppose, for example, that I am having a certain maritime experience; it will then be indeterminate whether my phenomenal properties include a hunk of wood or the ship it constitutes.

This way of setting the problem presumes a certain view about material constitution — namely, one according to which constitution is not identity, and there in the same location in the water there are two things, the ship and the quantity of wood which constitutes it. But this is not the only view in town, and if another response to the puzzles of material constitution — e.g. a four-dimensionalist one of the sort defended in Sider (2001) or a ‘dominant kinds’ view of the sort defended in Burke (1994) — were true, this problem would not arise. But a parallel puzzle would arise from permissive views about composition, according to which both the ship and the ship plus some piece of dust resting on the deck are genuine distinct objects; and in any case I would rather not pin my hopes on particular theories of material constitution.

Instead, I want to offer a reply in two parts: (i) this sort of indeterminacy is not immediately objectionable; and (ii) even if the trouble-making views about material constitution are true, I don’t see this kind of indeterminacy as an inevitable consequence of my view.

On (i): why would it be objectionable if it were sometimes indeterminate which phenomenal property you are instantiating? One of Pautz’s arguments against this way of going turns on the idea that there could be perceivers who can distinguish between phenomenal properties which differ only in which objects they involve — as mentioned in the preceding section, I don’t accept this possibility.

Another of his arguments is that if some indeterminacy is possible, then maybe a lot is — and I agree with him that an experience cannot represent an object as indeterminately red-or-green. It would be nice to have some explanation of why indeterminacy of one kind but not the other is possible. But this is hardly the only case where indeterminacy of one sort but not of another sort is possible. It can be indeterminate whether something is a tadpole or a frog; it cannot be indeterminate whether something is a tadpole or an asteroid. So some indeterminacy does not inexorably lead to unrestrained indeterminacy.

On (ii): suppose that we want to avoid any indeterminacy of phenomenal properties, and suppose to generate the difficulty that the relevant view of material constitution is correct. Then I must say that either the ship or the hunk of wood, but not both, is a constituent of the phenomenal property. Which one? Here is a flat-footed answer: I don’t know; this is a topic about which we might go on to theorize.

Pautz seems to assume that this sort of flat-footed answer is a non-starter. And I suspect that lurking in the background is the thought that my view faces a version of the ‘Benacerraf problem’ — surely both can’t be parts of the content; but nothing could make one but not the other a part of it; so neither is.

I fear that I am congenitally unable to feel the force of the Benacerraf prob-

lem, and believe that it is much too glibly deployed in contemporary philosophy. In my view, there is often a much too quick inference from ‘I can’t now see what could make X rather than Y the case’ to ‘Nothing could make X rather than Y the case.’ This is a perilous inference, not to be undertaken lightly. In the present case I feel pretty comfortable saying: if the trouble-making view about constitution is true, then something makes one rather than the other part of the content — I don’t know offhand what would, but that’s hardly a reason to think that nothing could.¹⁴

5 WORRIES ABOUT AVAILABILITY

In the previous section I tried to deflect various objections to my view that phenomenal properties include relations to objects and natural kinds. But Pautz also casts doubt on the central argument I give for that view, which turns on principles connecting the contents of experiences with the thoughts made available for thought by those experiences to subjects.

Let *generalism* be the view that sensing properties never include objects or kind properties. The best version of this response — the one that Pautz presses — concedes that experiences sometimes make objects and kinds available for thought, but doubts that this entails that the contents of the experiences themselves include objects and kinds.

One way to defend generalism is to look at cases in which experiences make objects available for thought without those objects being plausible candidates to be parts of the contents of experience. Pautz’s footprint example is just such a case. One experience makes available thoughts about Joe, and the other makes available thoughts about Bob; so, by Availability/Difference, there must be some difference in the contents of the experiences. But it is obvious that neither of the experiences represents Joe or Bob, and it is hard to see what other difference in content there could be.¹⁵ So, Availability/Difference is false.

Here is what I would have said to this sort of case when I was writing the book: the thoughts about Joe and Bob are not singular thoughts about these individuals, but rather descriptive thoughts of the form ‘[the x : x made this footprint] ... x ...’ So, despite initial appearances, there is no difference in thoughts made available; counterexample resolved.

¹⁴Pautz argues that I am committed to thinking that there will be indeterminacy of phenomenal properties at least in cases of ‘slow switching’ from one environment to another. But while this is a plausible view about the content of thoughts, it seems much less obvious to me in the case of experience. Suppose that there are identical twins, who confusingly both go by the name of ‘Bob.’ Suppose that I know only Bob-1, but that at a certain point in time Bob-1 moves far away, and is replaced (unbeknownst to me) by Bob-2. It is plausible that there is a period in which it is indeterminate whether ‘Bob’ out of my mouth refers to Bob-1 or Bob-2. But this is not an especially plausible view about the contents of my experiences; as soon as Bob-2 appears on the scene, I think that I am (determinately) visually representing him.

¹⁵One might say: one represents one footprint, and the other experience represents a numerically distinct footprint. That is a reasonable enough thing to say about this case, but it is easy to construct parallel cases which don’t permit this escape route.

I thought this because I held what Hawthorne and Manley (2012) call an ‘acquaintance theory’ of singular thought: the view that one can only have a singular thought about some object if one stands in some special relation to that object, one in which the subject of experience fails to stand to Joe and Bob in the above example. I’m now persuaded by the arguments of Hawthorne and Manley this sort of view is false.

I now think that cases of the sort Pautz discusses show that we need some sort of restriction on Availability/Difference. That is, we need some way of filling out the following principle:

If two experiences make available different thoughts of type X, then they must differ in content.

The obvious question is then: what is type X?

In the footprint case, here is what seems to be going on. First, the experience immediately makes available some objects and properties for thought — saliently, in this case, the footprint. The subject of the experience is then able to use these objects and properties to formulate new singular thoughts about other objects and properties — like the maker of the footprint. So we seem to have at least a rough distinction between the objects and properties which an experience makes directly or immediately available for thought, and those it makes indirectly available for thought. The natural suggestion is then to restrict the Availability/Difference principle to contents made immediately or directly available for thought.

In a sense, this is a version of my old view. The background thought is that the experience directly makes available certain objects or properties for thought, and one can then use those to formulate demonstrative thoughts where those directly-available contents are part of the restrictor. I used to say that these were not genuinely singular, but instead were disguised descriptions; the present suggestion gives this up, but instead classifies these as indirectly rather than directly made available by the experience.

Demonstrative thoughts about objects seem pretty clearly directly made available; and, while I concede that matters are less clear here, it is at least arguable that the relevant sorts of thoughts about natural kinds are as well.

It must be said that the history of spelling out distinctions like the one between contents made directly and those made indirectly available does not exactly inspire confidence. But the above gloss on the footprint case (and other parallel cases) does seem to be a natural one; it is hard to believe that the subject’s ability to have thoughts about the footprint does not have some explanatory priority over her ability to have thoughts about its maker.

In sum: for any thought made available by an experience, we can explain how it is made available by that experience. Either it is made available directly by the experience, or indirectly. In the latter case we can ultimately explain its availability in terms of objects and properties directly made available. In the former case we explain its availability by inclusion in the content of the

experience.¹⁶

REFERENCES

- Block, N. (1999). Sexism, racism, ageism and the nature of consciousness. *Philosophical Topics* 26(1-2).
- Burke, M. B. (1994). Preserving the principle of one object to a place: A novel account of the relations among objects, sorts, sortals, and persistence conditions. *Philosophy and Phenomenological Research* 54(3), 591–624.
- Byrne, A. and D. Hilbert (2007). Truest blue. *Analysis* 67, 87–92.
- Fine, K. (2007). *Semantic Relationism*. Blackwell Publishing.
- Hawthorne, J. and D. Manley (2012). *The Reference Book*. Oxford University Press.
- Hilbert, D. (2005). Color constancy and the complexity of color. *Philosophical Topics* 33(1), 141–158.
- Johnston, M. (2004). The obscure object of hallucination. *Philosophical Studies* 120, 113–183.
- Kripke, S. (1979). A puzzle about belief. In A. Margalit (Ed.), *Meaning and Use*, pp. 239–283. Boston: D. Reidel Publishing Company.
- Salmon, N. (1990). A Millian heir rejects the wages of *Sinn*. In C. A. Anderson and J. Owens (Eds.), *Propositional Attitudes: The Role of Content in Logic, Language, and Mind*, pp. 117–139. Stanford, CA: CSLI Press.
- Salmon, N. (2012). Recurrence. *Philosophical Studies* 159(3), 407–441.
- Sider, T. (2001). *Four Dimensionalism: An Ontology of Persistence and Time*. Oxford University Press.
- Soames, S. (2015). *Rethinking Language, Mind, and Meaning*. Princeton University Press.

¹⁶Thanks to Casey O’Callaghan for very helpful comments on a previous version of this reply.