

Spectrum inversion without a difference in representation is impossible

Jeff Speaks

Published online: 29 August 2010
© Springer Science+Business Media B.V. 2010

Abstract Even if spectrum inversion of various sorts is possible, spectrum inversion without a difference in representation is not. So spectrum inversion does not pose a challenge for the intentionalist thesis that, necessarily, within a given sense modality, if two experiences are alike with respect to content, they are also alike with respect to their phenomenal character. On the contrary, reflection on variants of standard cases of spectrum inversion provides a strong argument for intentionalism. Depending on one's views about the possibility of spectrum inversion, the impossibility of spectrum inversion without a difference in representation can also be used as an argument against a variety of reductive theories of mental representation.

Keywords Spectrum inversion · Perception · Intentionalism · Representationalism

Questions about the possibility of spectrum inversion are questions about whether the phenomenal characters, or phenomenologies, of the experiences of two subjects can differ in certain ways even if various facts about them are held fixed. So we can ask, for example, whether the phenomenologies of the color experiences of two agents can systematically differ even if the two agents are alike with respect to their behavioral dispositions, their functional organization, or their physical constitution. So put, the possibility of such scenarios has obvious relevance to attempts to give an account of the phenomenal character of perceptual experience in terms of, for example, behavioral dispositions, functional organization, or physical constitution.

J. Speaks (✉)
University of Notre Dame, 100 Malloy Hall, Notre Dame, IN 46556, USA
e-mail: jspeaks@nd.edu

1 Spectrum inversion and intentionalism

But intuitions about the possibility of spectrum inversion can also be aimed at a target other than attempted reductions of phenomenology: they can, it seems, be used to refute any view which claims that there is a necessary connection between the *content* and phenomenology of perceptual experience.

Discussion of the content of experience, while widespread, has a less well-entrenched place in the philosophical literature than discussions of phenomenal character.¹ We can gloss ‘the content of an experience’ as the way that experience presents, or represents, the world as being; the way the world is, according to the experience; the way that the world would have to be for the experience to be veridical. One way to get a grip on the notion of the content of experience is via the phenomenon of illusions. Intuitively, an illusion is an instance of some sort of misrepresentation; so, plausibly, in illusions there must be some contentful state of the perceiving subject whose content is a false proposition. But this can’t be the state of the agent judging or believing some false proposition, since it’s possible for subjects to experience illusions without being at all convinced by them. A natural thought is that the contentful state in question is the perceptual experience itself; if this is right, then it makes sense to talk about the contents of perceptual experiences.

To be sure, not everyone agrees that perceptual experiences have contents, and there are treatments of illusory experience available which avoid commitment to the view that experiences have contents.² The aim of the foregoing isn’t to convince skeptics that experiences have contents, but just to go some distance toward making talk about the content of experience intelligible as something distinct from the contents of judgements, beliefs, or other more familiar propositional attitudes. In what follows, I’ll take for granted that experiences have both content and phenomenal character, and that it makes sense to talk about two experiences (whether of one or two subjects) being alike or different with respect to their contents, and alike or different with respect to their phenomenal characters.

The view about the relationship between content and phenomenal character that we’ll be concerned with is the following:

Interpersonal Intentionalism

Necessarily, if two experiences of any two subjects (of the same sense modality) differ in phenomenal character, then they differ in content.

It is important to distinguish this view from other views which go under the names ‘intentionalism’ and ‘representationalism.’ The above claim does not immediately entail that phenomenology is identical to a certain kind of content, that phenomenology is reducible to a certain kind of content, that two experiences have the same phenomenal character if *and only if* they have the same content, or that any mental state or event with a phenomenology also has representational content—though the

¹ At least within the analytic tradition; talk about perceptual states as intentional states has been a more or less constant feature of discussions of perception in the phenomenological tradition since Brentano.

² For examples of skepticism about the applicability of the notion of content to perceptual experience, see Chap. 6 of Campbell (2002), Travis (2004), Alston (2005), and Brewer (2006). For a discussion of illusion from the perspective of such a skeptic about perceptual content, see Brewer (2008).

first three of these views do entail Interpersonal Intentionalism. Nor, obviously, does Interpersonal Intentionalism, as formulated above, rule out any of these stronger theses. Interpersonal Intentionalism just says that, within any perceptual modality, phenomenology supervenes on content; it is reasonable to think that this is the minimal claim that someone should have to endorse to count as an intentionalist.

Later it will be important to distinguish Interpersonal Intentionalism from other, weaker supervenience theses, which restrict the supervenience of phenomenal character on content to the experiences of a single subject, or to the experiences of a single subject within a restricted time interval. In the mean time, I will just use ‘intentionalism’ as a shorthand for Interpersonal Intentionalism, as formulated above.

It is important to see that the fact that the thesis stated above is weaker than the views which often go by the names ‘intentionalism’ and ‘representationalism’ does not make it uncontroversial; far from it. The philosophy of mind is divided between those philosophers who endorse this thesis, and those who think that phenomenology can ‘outrun’ representational content, in the sense that two experiences can differ in phenomenology without an accompanying difference at the level of content.³ Perhaps the strongest and most influential motivation for the latter position goes by way of the possibility of the kind of spectrum inversion which was introduced by John Locke:

Neither would it carry any Imputation of Falsehood to our simple Ideas, if by the different Structure of our Organs, it were so ordered, That the same Object should produce in several Men’s Minds different Ideas at the same time; v.g. if the Idea, that a Violet produced in one Man’s Mind by his Eyes, were the same that a Marigold produces in another Man’s, and vice versa. For since this could never be known because one Man’s Mind could not pass into another Man’s Body, to perceive, what Appearances were produced by those Organs; neither the Ideas hereby, nor the Names, would be at all confounded, or any Falsehood be in either. For all Things, that had the Texture of a Violet, producing constantly the Idea, which he called Blue, and those which had the Texture of a Marigold, producing constantly the Idea, which he as constantly called Yellow, whatever those Appearances were in his Mind; he would be able as regularly to distinguish Things for his Use by those Appearances, and understand, and signify those distinctions, marked by the Names Blue and Yellow, as if the Appearances, or Ideas in his Mind, received from those two Flowers, were exactly the same, with the Ideas in other Men’s Minds. (*Essay on Human Understanding*, Sect. II.xxxii.15)

If Locke’s scenario is possible, the argument against intentionalism seems straightforward. Locke supposes that the same object might “produce in several Men’s Minds different Ideas at the same time”—i.e., the same object might produce in several perceivers experiences with different phenomenal characters—even though there would be no “Falsehood ... in either.” But if there is no falsehood in either’s representation of the color, then—given that the object has just one color—they must

³ See, e.g., the statement of ‘phenomenism’ in Block (1995). For other arguments for anti-intentionalism, see Peacocke (1983), Boghossian and Velleman (1989), Block (1990), and Macpherson (2005).

be representing it as the same color. But then, supposing that there is no difference in the contents of the two experiences aside from their representation of color, we have a difference in phenomenology without a difference in content, and Interpersonal Intentionalism must be false.

One assumption of the Lockean argument against intentionalism which many contemporary color theorists would reject is the assumption that the relevant object has ‘just one color’, rather than being blue-to-one-observer-in-his-circumstance and yellow-to-the-other-observer-in-her-circumstance. If this sort of relativist view of color is correct, then the argument from inversion without misrepresentation against intentionalism is simply defused, because the inverted subjects can differ in their representation of the color of the relevant object without either misrepresenting it.⁴ In what follows, I’ll be setting this sort of view of color to the side, and asking whether, given the assumption that color relativism is false, there can be spectrum inversion without a difference in representation. As this amounts to simply granting a key premise to the Lockean anti-intentionalist, this assumption can hardly be construed as stacking the deck in favor of intentionalism.

Though I will return to this topic below, for now I am also setting to the side the attempt to make Interpersonal Intentionalism consistent with spectrum inversion without misrepresentation by explaining the phenomenal differences between inverted perceivers via differences in their representation of phenomenal (or appearance) properties, rather than differences in their representation of color properties.⁵ As with color relativism, ignoring this response to the Lockean argument gives, if anything, the advantage to the anti-intentionalist. I’ll be asking whether spectrum inversion without misrepresentation is possible given, for the purposes of argument, the dual assumptions that color relativism is false and that there are no properties with respect to which the contents of the inverted experiences could differ other than color properties. Given these two assumptions, the claim that spectrum inversion without a difference in representation is possible is equivalent to the claim that spectrum inversion without misrepresentation is possible.

Why should we acquiesce in Locke’s claim that neither of the subjects spectrum-inverted relative to each other is misperceiving the color of the relevant object? The intuition behind that claim is not hard to generate:

It is hard to take seriously the idea that we are systematically misrepresenting the colors of objects. But subjects phenomenally inverted relative to us are really, from an epistemic point of view, *just like us*. It would be just as hard for them to take seriously the idea that they are systematically misrepresenting the colors of objects. And for good reason—their judgements about colors correlate well with those of their companions, and they have proven to be a very reliable guide in making their way about the world. So the intentionalist has two options, neither appealing. Either he can claim, ludicrously, that *our* color experiences are systematically misrepresenting the colors of things, or he

⁴ For defenses of views of color of this sort, see Johnston (1992), McLaughlin (2003), Cohen (2003, 2004, 2007). A related view which also blocks the Lockean argument is the color pluralism of Kalderon (2007).

⁵ For a defense, see, among other places, Shoemaker (2001).

can endorse an unreasoned chauvinism which convicts the occupants of Inverted Earth of the same systematic misrepresentation.

This sort of reasoning can make the view that spectrum inversion without misrepresentation is possible seem less theoretically loaded than the view that spectrum inversion simpliciter is possible among behavioral, functional, or physical duplicates. For what we need is just that two subjects—who may differ physically, functionally, and behaviorally—can be spectrum inverted relative to each other, even though it seems implausible to attribute systematic error to the perceptual experiences of either.

There are various ways of turning intuitions which favor the possibility of spectrum inversion without misrepresentation into an explicit argument against intentionalism; and intentionalists have for the most part responded to these arguments by finding premises in them to reject.⁶ My aim in this paper is different: to give a *reductio* of the idea that Lockean examples of spectrum inversion without misrepresentation are possible. Afterwards, I'll return to the consequences of this result for the truth of intentionalism, and for theories of mental content.

2 Radical anti-intentionalism

Suppose that spectrum inversion without misrepresentation is possible, and two experiences can be alike at the level of content while differing in phenomenal character. In that case, what *is* the connection between phenomenology and content? In particular, we can ask: is it a necessary truth that there be *some* systematic correlation, for any given subject, between phenomenology and content?

Suppose first that the proponent of spectrum inversion without misrepresentation says 'No'. This yields the following *radical anti-intentionalist* view:

Content and phenomenal character are independent aspects of perceptual experience. It so happens that in our case there is a systematic (though contingent) connection between experiences with a certain phenomenal character ('red-feeling' experiences, for example) and experiences which represent objects as having certain color properties (redness, for example). But that is just an accidental feature of our constitution and relation to our environment; there could have been creatures for whom there was no systematic connection between content and phenomenology.⁷

Radical anti-intentionalism is very implausible. For, if radical anti-intentionalism were true, then the following scenario would be possible:

A subject is looking intently at a well-lit surface which occupies the whole of the subject's visual field. Over the course of a few seconds, his experience

⁶ For clear discussions, see Hilbert and Kalderon (2000) and Marcus (2006).

⁷ The use of 'red-feeling' as a term for 'having the phenomenology characteristic of my experiences which represent an object as red' is borrowed from Byrne and Hilbert's introduction to Byrne and Hilbert (1997).

goes from being (as we would put it) bright-red-feeling to being bright-green-feeling to being bright-blue feeling, and constantly repeats this pattern. But, the whole time, he is visually representing the wall as yellow; it visually seems to him throughout that the wall is yellow; according to his experience, the wall is yellow throughout.

In asking whether this scenario is possible, it is important to be clear that the relevant question is *not* whether an agent could have an experience with this rapidly changing phenomenology while ‘representing the color of the wall as yellow.’ Of course he could. The agent could be, throughout, thinking that the wall is yellow, judging that the wall is yellow, and believing that the wall is yellow; and these are all ways of generically representing that the wall is yellow. The question is whether an agent could have an experience with this kind of changing phenomenology while *visually* representing the color of the wall as a constant shade of yellow. Once the question is put this way, it is clear, it seems to me, that the above does not give a description of a possible situation; if your visual experience is by turns red-feeling, green-feeling, and blue-feeling, you are not visually representing as constant the color of the object to which you are attending. We can put this by saying that it is impossible to combine psychedelic color phenomenology with constant color content.

The same point can be made, though in a less phenomenologically vivid way, by trying to imagine the inverse scenario:

A subject is looking intently at a well-lit surface which occupies the whole of the subject’s visual field. The only thing notable about the phenomenology of his experience of the surface is its monotony. The experience is charcoal-gray-feeling, and remains so for the duration of the experience. Nonetheless, the subject is visually representing the color of the wall as constantly changing from bright red, to bright green, to bright blue; it visually seems to him that the wall is changing from bright red, to bright green, to bright blue; according to his experience, the wall is changing from bright red, to bright green, to bright blue.

Let’s say that a series of experiences has psychedelic *content* if it represents the color of the relevant surface as exhibiting dramatic changes in color. The incoherence of the scenario just described indicates the impossibility of combining constant phenomenology with psychedelic content.

Most people to whom I have presented these scenarios have agreed that they are impossible, and even obviously so. But because the judgement that these scenarios are impossible will be essential to the argument which follows, it is worth considering some possible objections.

First, one might object to the idea that cases of psychedelic color phenomenology + constant color content are impossible by drawing attention to the phenomenon of color constancy. Cases of color constancy are cases in which the phenomenal character of our visual experience of a surface changes (with, for example, changes in lighting) without any change in our representation of the color of the relevant surface. But if cases of this sort are possible, why not cases of psychedelic color phenomenology along with constant color content?

A first point to note about this response is that cases of color constancy are only a *prima facie* challenge to the impossibility of psychedelic phenomenology + constant content; they say nothing to the claim that instances of constant phenomenology + psychedelic content are impossible. But in the end it is doubtful whether these cases are a serious challenge to either of the claims about impossibility given above.

On one plausible view of the phenomenon, in cases of color constancy, the human visual system responds to certain changes in the environment by representing one property—the color of the surface—as fixed, and representing another property—the degree to which the surface is illuminated—as changing.⁸ But the case described above is one in which all aspects of the phenomenology of the subject's experience other than color phenomenology are held fixed—in particular, we hold fixed the degree to which the experience is (as we might put it) light-feeling or dark-feeling. (One way of making clear the irrelevance of illumination to the above cases is by imagining variants of the above cases in which the subject is looking directly at light sources which seem to be changing colors rather than at opaque surfaces.) So if the above view of color constancy is correct, it does not seem as though we can interpret examples of psychedelic phenomenology as of a piece with examples of color constancy.⁹

But the impossibility of psychedelic phenomenology + constant content is not hostage to views of color constancy which understand them as cases in which represented color and represented illumination come apart.¹⁰ Let's suppose that, despite our stipulation that there are no changes in illumination, it is possible for a scenario like that described under the heading of psychedelic phenomenology + constant content to be one in which the wall is represented as having a constant color while being placed under differently colored lights. Still, on this interpretation, it is surely true that *something* is being represented as changing color; whether it is the wall, or the light under which it is being viewed, is irrelevant to what follows.¹¹

A quite different sort of objection to the intuition that the above cases are impossible is that this intuition is based on failure to adequately separate two senses of 'seems.'¹² It might seem implausible to say that (in the case of psychedelic phenomenology + constant content) the color of the wall could *visually seem* to be yellow throughout—but when we have this intuition we are understanding 'visually seems' in a phenomenal rather than an epistemic sense, and of course it is part of the description of the case that the phenomenology of the experience is changing rather

⁸ See, for example, Hilbert (2005).

⁹ Similar remarks are in order about the 'same color illusion' and other cases in which the representation of the color of a surface appears to differ depending on the colors of surrounding surfaces. We can simply abstract from contextual dependence of this sort by imagining that, in the above cases, the experience of the colored surface occupies the whole of the subject's visual field.

¹⁰ For an alternative to views of this sort, see Cohen (2008).

¹¹ Thanks to an anonymous referee for helpful discussion of the relevance of color constancy to the above argument.

¹² For discussion of the distinction between epistemic and phenomenal senses of 'seems,' see, among other places, Chisholm (1957).

than constant. But this hardly shows that the color of the surface visually seems to be changing if we interpret ‘seems’ in its epistemic sense—after all, I might well come to know that when confronted with a yellow surface, my visual system responds with a series of experiences which exhibit just the psychedelic phenomenology described above. In such a case, my visual experience would provide evidence that the surface was a constant yellow—so, in the epistemic sense of ‘seems’, it is not impossible for the visual experience to be one in which the surface of the wall seems to be a constant yellow.¹³

This point about the epistemic sense of ‘seems’ is correct as far as it goes, but irrelevant to the question of whether cases of psychedelic phenomenology + constant content are possible. Of course the subject could, through the series of experiences which exhibit psychedelic phenomenology, judge that the color of the wall is a constant yellow; furthermore, the subject’s experiences could, given the background beliefs about her visual system described above, provide her with evidence for this judgement. But what is in question is not what it is possible for the contents of the subject’s judgements or beliefs to be, but what it is possible for the contents of the subject’s visual experiences to be. One might, of course, deny that there is such a thing as the content of a subject’s visual experience, as distinct from the contents of the subject’s judgements or beliefs; but if we grant that there is such a thing as the content of visual experience, there is no easy inference from a claim about the contents of a subject’s judgements on the basis of experience to a claim about the contents of that experience. After all, in cases in which a rational subject knows that she is having an illusory experience, we can expect the two to diverge.¹⁴

3 Moderate anti-intentionalism

On the basis of the impossibility of the scenarios described above, the proponent of the possibility of spectrum inversion without misrepresentation should recoil from radical anti-intentionalism to *moderate anti-intentionalism*: she should say that, even though Lockean examples show that there is no necessary connection between the phenomenal characters and contents of two arbitrarily chosen experiences, it is yet a necessary truth that there be, for any subject, *some* systematic connection between phenomenology and content.

There is nothing initially incoherent in the idea that phenomenology and content might be related in this way. A useful analogy here might be the relationship between linguistic expressions and what those expressions signify. There is clearly no necessary connection between an expression and the object for which it stands; but it is plausibly a necessary truth that, whenever a linguistic expression stands for

¹³ One might give a similar objection to the idea that cases of constant phenomenology + psychedelic content are impossible, based on a case in which you know that your visual system generates a constant charcoal-gray-feeling experience only in response to surfaces which are rapidly changing from bright red, to bright green, to bright blue.

¹⁴ A different sort of objection than the two considered is that in the case of psychedelic phenomenology, what must be changing is the representation of ‘appearance properties’, not color properties. I’m ignoring this possibility for the moment—I return to it below.

an object, there must be some connection between them which could explain how the one came to stand for the other. So perhaps color-phenomenology and color-content are related in something like the way that expressions and their referents are related; for any given agent, experiences with a certain color-phenomenology *represent* the color by which experiences of that phenomenal type are typically caused.

This model of the relationship between phenomenology and content has the virtue of fitting nicely with the Lockean intuitions sketched above: two subjects spectrum inverted relative to each other can represent the same color property with phenomenally different experiences because, for each of them, experiences with that phenomenology are correlated with the relevant color. Imagine two subjects whose vocabulary was inverted in this way: one uses 'green' in a situation just in case the other would use 'red' in that situation, and uses 'red' in a situation just in case the other uses 'green' in that situation.¹⁵ We wouldn't hesitate to describe this as a case of 'vocabulary inversion without misrepresentation': one will predicate 'red' of an object just in case the other predicates 'green' of that object, yet neither is systematically misattributing colors to objects, since 'green' out of one's mouth predicates the same property as 'red' out of the other's mouth. Shouldn't we say that green-feeling experiences stand to color properties in something like the same relation in which color words stand to color properties?

One worry about this analogy emerges if we note that linguistic representation permits both coreferential terms and ambiguity: there is nothing unusual or controversial about the idea that two genuinely distinct linguistic symbols can stand for the same object or property, or that one symbol can stand, on different occasions, for two distinct objects or properties. But a thought experiment much like the one used against the radical anti-intentionalist can be used to show that this aspect of the relationship between symbols and their referents can't be carried over to the relationship between experiences with a certain color-phenomenology and color-properties. For suppose that 'coreference' of this sort were possible; suppose, to fix ideas, that for a given agent green-feeling and purple-feeling experiences represent the same color property. Then the agent could have an experience which flipped rapidly back and forth between (as we would put it) a green-feeling phenomenology and a purple-feeling phenomenology, while nonetheless visually representing the color of the relevant surface as constant. But this is incoherent, in the same way that the above example of psychedelic phenomenology + constant content is incoherent. It is impossible for an agent's experience to be constantly switching between green-feeling and purple-feeling, while the whole time it visually seems to the agent that the surface in question has a constant color. We could run a similar argument using the possibility of 'ambiguity'; in this case, the counterexample would be a version of constant phenomenology + psychedelic content.

The moderate anti-intentionalist is likely to complain that this rests on taking the analogy with linguistic expressions and their referents too seriously. (Perhaps she

¹⁵ Here I'm thinking of the two language users as alike at the level of perceptual content and phenomenology—not as spectrum inverted as well as vocabulary inverted. This is important for considering certain kinds of inverted spectrum arguments against intentionalism. See the discussion of Thau in note 33.

will claim that the possibility of coreference and ambiguity in the case of linguistic expressions is explained by their being conventional; this would then be a disanalogy with the case of perceptual representation.) The connection between phenomenology and content, the moderate anti-intentionalist should say, is like the relationship between expressions and their referents in some ways, but unlike it in others; and one of the ways in which it differs is that, as a matter of necessity, it is never the case that, for any one agent (or for any one agent in a suitably short interval), experiences of distinct phenomenal types represent a single color, or that experiences of a single phenomenal type can on different occasions represent distinct color properties.

One way to put this is by understanding the position of the moderate anti-intentionalist as an instance of the following claim:

For any subject, there is some relation R which is such that experiences with color-phenomenology F represent color-property G if the F -feeling experiences bear R to G .

The impossibility of psychedelic phenomenology + constant content and constant phenomenology + psychedelic content can then be thought of as putting a constraint on the relation R : R must be a one-to-one relation between experiences with certain types of phenomenal character and the color properties represented by those experiences.¹⁶

This constraint rules out several otherwise plausible candidates for R , such as that experiences with a given phenomenology represent a certain color property just in case those experiences indicate the property. Let's say that a state of a certain kind indicates x for a subject iff were optimal conditions to obtain, the subject in question would come to be in that state only because of x .¹⁷ Clearly, there's nothing to rule out two distinct states indicating the same thing; just imagine that, in optimal conditions, x sometimes causes one state and sometimes another, and that nothing else (in optimal conditions) ever causes either of those states. The same goes for teleological theories, which might try to explain the connection between content and phenomenology for a creature in terms of facts about what evolutionary role was played by perceptual states with a certain phenomenology in the history of the creature's species. There's clearly no impossibility in states with two distinct color phenomenologies both playing the same evolutionary role in the history of a species. Since neither of these make R a one-to-one relation, neither are plausible

¹⁶ We are not licensed to build in the assumption that the same relation R must do this work in the case every agent; just that any relation which plays this role must be a one-to-one relation between experiences of certain phenomenal types and the color properties they represent. Maybe, for all we have said, there are a number of irreducibly distinct ways in which a linguistic expression can come to have a given referent, and perhaps, just so, there are a number of irreducibly distinct ways in which a color-phenomenology can come to be associated (for a given agent) with a certain color property. We're also not licensed to assume that the same relation must do this work at every time in the life of a single agent. These points are irrelevant to the arguments of this section; I return to them explicitly in the discussion of Quietist Moderate Anti-Intentionalism in the section which follows.

¹⁷ See, among many other places, Stampe (1979) and Stalnaker (1984). One, could, of course, stipulate that indication is a one-to-one relation, by adding to the definition the requirement that x can only indicate y if nothing else does. This sort of view would be open to the objection to covariational theories discussed below.

candidates for explaining the representation relation between phenomenal character and color.

But other candidates for *R* might seem more promising, such as the view that experiences with a given phenomenology represent a certain color property just in case they (under certain conditions) *co-vary* with that property.¹⁸ Since a subject is able to have at most one visual experience with color phenomenology at a time, it is plausible to think that it is impossible for experiences of distinct phenomenal types to each, for a single subject, *co-vary* with a single color property. For this reason, it seems that identifying *R* with some sort of covariation correctly rules out examples of psychedelic phenomenology and constant content, as well as the converse examples of psychedelic content and constant phenomenology.

However, covariational accounts are open to the same form of objection as the accounts considered above: they entail that a clearly impossible situation is possible. Consider the following scenario:

Full color phenomenology + no color content

A subject is (still) looking intently at a well-lit surface. The subject's experience has the same phenomenology as your experience of something bright red. However, the subject's perceptual experience does not represent the color of the surface at all; in fact, it does not represent anything as having any color. So, when the subject's experience changes its phenomenal character to one similar to the phenomenal characters of the experiences you have when looking at bright green things, the subject's experience does not represent the color of the wall as having changed.

Like the examples of psychedelic phenomenology + constant content and constant phenomenology + psychedelic content, the case of full color phenomenology + no color content is, it seems to me, clearly impossible. Indeed, this is just a special, particularly absurd case of psychedelic phenomenology + constant content.¹⁹

But, if *R* were some sort of covariation, it would be possible. Suppose that when a creature's visual system comes across a green object (under the right conditions), it sometimes triggers experiences with a green phenomenology, and sometimes experiences with a red phenomenology. Then neither type of experience—neither those with red phenomenology, nor those with green phenomenology—will covary with either color property. Since green surfaces will sometimes trigger experiences with red phenomenology, experiences with green phenomenology can't covary with greenness; since green surfaces will sometimes trigger experiences with green phenomenology, experiences with red phenomenology won't covary with greenness. So neither type of experience—neither those with red phenomenology nor those with green phenomenology—will covary with green, and neither will covary with instantiations of any other color. So, in this sort of case, the covariational theory yields the result that when the agent in question has experiences of this sort,

¹⁸ See, for example, Tye (1995, 2000).

¹⁹ It's important to keep in mind here that we're talking only about perceptual experiences, and not about 'phenomenal states' more generally. So I'm not here denying Block's plausible claim that, e.g., 'orgasm experiences' lack a representational content. See Block (2003).

she simply fails to represent the objects in question as having a color; it does not visually seem to the agent that the objects in question have a color. But this is extremely implausible.

If it does not seem immediately obvious that this is an absurd result, it may help to consider ‘mixed’ cases. We can imagine that experiences of the agent which have a yellow phenomenology or an orange phenomenology, do satisfy the requirements of the covariational theory even though experiences with a red phenomenology, as above, fail to covary with any color property. But in this kind of case, if the covariational theory is to be believed, if the subject is looking at a screen with colors being projected upon it, and her experience switches from yellow phenomenology to orange phenomenology to red phenomenology, what has happened is that the screen first visually seemed yellow to the subject, then visually seemed orange to the subject, and then ceased to seem to have any color at all. This is hard to believe. Surely the switch from orange phenomenology to red phenomenology can’t be a switch from representing the relevant surface as having a color to simply failing to do so. In general, covariational accounts make it more difficult to represent colors than it plausibly is.²⁰

So far we’ve focused on attempts to explain *R* in terms of ‘vertical’ relations between color experiences and the external features which, under various conditions, cause the subject in question to have color experiences of the relevant type. Perhaps this was our mistake. Why not think that the moderate anti-intentionalist should explain *R* in terms of ‘horizontal’ relations between experiences with the relevant sort of phenomenal characters? The analogy here is with the distinction between broadly causal and conceptual role theories of the contents of mental states, where the former explain the contents of symbols in terms of symbol-world relations, and the latter explain the contents of symbols in terms of the relations between internal states. A ‘conceptual role’ approach to the relationship between phenomenology and content seems well-suited to avoid the problems we’ve been discussing, since a theory which explains content in terms of relations between phenomenal states needn’t allow distinct phenomenal states to represent the same color property.

However, it’s clear that the moderate anti-intentionalist can’t make do with these sorts of relations alone. After all, if color content is explained wholly in terms of relations between certain phenomenal states, then color content should supervene on phenomenology. But this is just what the proponent of the possibility of spectrum inversion without misrepresentation must deny. To see this, imagine a pair of red-green spectrum inverted subjects, one of which is looking at a red apple and one of

²⁰ I think that the best reply for the covariational theorist is to deny the possibility of the sorts of cases described above: to deny that it is possible for experiences of a certain phenomenal type to fail to covary with color properties in the conditions specified by the theory. Another way to put this is to say that the conditions specified by the theory are such that it is impossible for experiences of the relevant phenomenal types to fail to covary with colors when those conditions obtain. But this is hard to believe. The ‘conditions specified by the theory’ must be loose enough to allow for spectrum inversion, so they must be loose enough to allow my phenomenal-green experiences to covary with a different color than do those of my invert. But how could they be loose enough to allow this without also allowing for the possibility that someone’s phenomenal-green experiences could fail to covary with any color property?

which is looking at a green apple. Since the subjects are spectrum-inverted, their experiences will presumably have the *same* phenomenal character in this case; but the proponent of the possibility of spectrum inversion without misrepresentation will want to allow the possibility that each represents the color of the apple they're viewing correctly. Since the apples differ in color, this would be a pair of experiences with the same color phenomenology but different color content; since the proponent of spectrum inversion has to recognize the possibility of such a pair of experiences, she has to deny that supervenience of color content on color phenomenology.

The natural move is then for the moderate anti-intentionalist to shift to a theory which explains *R* in terms of some combination of vertical and horizontal relations.²¹ The idea would then be that the horizontal relations could secure distinctness of content between distinct phenomenologies for an individual at a time, while the vertical relations could secure variance of content between individuals (or between distinct times in the life of a single individual).

But the problem is that it's hard to see how these two sorts of relations could be combined without yielding the problems faced by covariational accounts. Horizontal relations between experiences of certain phenomenal types are introduced to avoid the problem of experiences of distinct phenomenal types representing the same color property; we can therefore think of the use of these horizontal relations as imposing the requirement that distinct phenomenal types represent distinct colors.²² But now suppose that, for a subject, experiences of distinct phenomenal types P_1 and P_2 bear the relevant vertical relations to a single color property *C*. By parity, experiences of type P_1 represent *C* if and only if experiences of type P_2 do. But if both do, then scenarios of psychedelic phenomenology + constant content are possible; the subject's experience could switch rapidly between distinct phenomenal characters P_1 and P_2 while representing the color of the relevant surface as constant. But cases of this sort are impossible, so both can't represent *C*; so neither can. But presumably neither represents any color other than *C*. So experiences with phenomenal characters P_1 and P_2 do not, for this subject, represent objects as having any colors at all. But this is just a case of full color phenomenology + no color content, and, as argued above, cases of this sort are not possible. So combinations of vertical and horizontal relations fare no better than either taken singly.

The cases discussed above can be thought of as constituting a kind of dilemma for the moderate anti-intentionalist: (i) if a candidate for *R* is tolerant enough to rule out cases of full color phenomenology + no color content, then *R* will fail to be one-to-one, and so will make possible either or both of the scenarios of psychedelic

²¹ This is the analogue of the move from solipsistic to non-solipsistic conceptual role semantics, in the terminology of Harman (1987).

²² These horizontal relations might impose further requirements to do with relative location in color space, or subsumption of certain similar determinate shades under single determinate colors, as in the theory of Hilbert and Kalderon (2000). Though Hilbert and Kalderon are intentionalists, it seems to me that an argument of the present sort might be used against their theory of the contents of color experiences. In future work, I hope to show how arguments of the present sort might be brought to bear on functionalist theories of the content of color experience more generally, whether those are accompanied by intentionalist or anti-intentionalist views of the relationship between phenomenal character and perceptual content.

phenomenology + constant content or psychedelic content + constant phenomenology; but (ii) if a candidate for R is made demanding enough to block the possibility of these cases, it will end up making color representation implausibly difficult, making cases of full color phenomenology + no color content possible.

Perhaps there is some candidate for R which I have overlooked, and which would avoid both horns of this dilemma while still making spectrum inversion without misrepresentation possible. But this seems to me unlikely. We've considered theories which try to explain the relationship between phenomenology and content in terms of vertical relations between phenomenal states and the world, horizontal relations between phenomenal states, and combinations of the two. In each case, we've seen that the theories face either problem (i) or problem (ii), and a pattern has emerged: refinements to the theory which avoid the relevant horn of the dilemma end up pushing the theory onto the other. At this stage, it seems reasonable to think that there is no candidate for relation R which will do the work that the moderate anti-intentionalist wants it to do.

4 Quietist moderate anti-intentionalism

Here, though, the anti-intentionalist may wish to raise a question about the dialectical situation. Let us suppose that we have failed to specify a plausible candidate for the relation R . Why should the moderate anti-intentionalist feel pressed to specify this relation? The moderate anti-intentionalist will want to distinguish between the following three theses:

Intrapersonal Time-Restricted Intentionalism

Necessarily, if two experiences of a single subject within some minimal time interval t differ in phenomenal character, then they differ in content.

Intrapersonal Time-Unrestricted Intentionalism

Necessarily, if two experiences of a single subject (whatever the interval between them) differ in phenomenal character, then they differ in content.

Interpersonal Intentionalism

Necessarily, if two experiences (whether of a single subject or two subjects) differ in phenomenal character, then they differ in content.

The radical anti-intentionalist denies each of these; but the moderate anti-intentionalist is committed only to denying the third. The foregoing argument against moderate anti-intentionalism amounts, basically, to a demand that the moderate anti-intentionalist specify that relation between phenomenology and content which will make Interpersonal Intentionalism false, while letting at least one of the two versions of Intrapersonal Intentionalism be true. But perhaps the moderate anti-intentionalist can resist this demand for explanation. The two versions of Intrapersonal Intentionalism are just supervenience theses, as is Interpersonal Intentionalism. Why can't the moderate anti-intentionalist simply accept the first or first and second of these theses, deny the third, and be done with it?

Like most assertions of supervenience claims without explanation of the supervenience, there's something unsatisfying about this sort of quietist moderate anti-intentionalism.²³ But can we give any argument against it? I think so.

First, suppose that the quietist moderate anti-intentionalist endorses only Intrapersonal Time-Restricted Intentionalism, and rejects Intrapersonal Time-Unrestricted Intentionalism. In this case, if two consecutive experiences of a subject have the same color content, they must have the same color phenomenology; but if the two experiences are separated by some minimal interval of time—call this t —they might have the same content, and different phenomenology.²⁴ But presumably it is possible for the subject to be having perceptual experiences during t , which have some color phenomenology. Since, by hypothesis, t is the minimal interval of time by which two experiences alike in color content but distinct in color phenomenology must be separated, they cannot have the color content appropriate to the perceiver's situation prior to t ; but by the same reason they can't have the color content appropriate to the perceiver's situation subsequent to t . And they can't have some third sort of content since, as in the previous cases, this would violate the stipulation that t is the minimal interval of time by which two experiences alike in color content but distinct in color phenomenology must be separated. So they must have no color content. But this would, contra our conclusions above, make cases of full color phenomenology + no color content possible. To avoid this interval of contentless perceptual experience, it must be the case that $t = 0$. But this is to collapse our first version of quietist moderate anti-intentionalism to radical anti-intentionalism; if $t = 0$, then psychedelic phenomenology + constant content would be possible. But it isn't. So the anti-intentionalist should not reject Intrapersonal Time-Unrestricted Intentionalism unless she wants also to reject Intrapersonal Time-Restricted Intentionalism, and so become a radical anti-intentionalist.²⁵

So suppose that the quietist moderate anti-intentionalist endorses Intrapersonal Time-Unrestricted Intentionalism but rejects (as he must, to qualify as an anti-intentionalist at all) Interpersonal Intentionalism. Let A and B be distinct individuals whose experiences, by the anti-intentionalist's lights, differ in phenomenology but have the same content. Whatever the facts about these individuals are which determine the phenomenologies and contents of their mental states, can't we always imagine a sufficiently long-lived individual who, at one stage in his life, is identical

²³ Of course, the same complaint might be made against the Interpersonal Intentionalist who fails to offer any explanation of this supervenience claim.

²⁴ Here I'm again assuming that color phenomenology is linked to representation of color properties, and setting to the side views which make use of the representation of appearance properties. I discuss these views below.

²⁵ There is a kind of analogy here to sorites arguments and the range of cases in which, intuitively, it is indeterminate whether a predicate applies to a thing. Truth-value gap approaches are the analogue of the interval of contentless experience, and epistemicism is the view that $t = 0$. One thought is that one or another view about vagueness might come to the aid of the Intrapersonal time-restricted intentionalist who rejects Intrapersonal Time-Unrestricted Intentionalism. So far, I haven't been able to come up with any plausible candidates.

in all relevant respects to *A*, and at another stage, is identical in all relevant respects to *B*? If so, then it is hard to endorse Intrapersonal Time-Unrestricted Intentionalism without endorsing Interpersonal Intentionalism.²⁶

Just as Intrapersonal Time-Restricted Intentionalism (plus the denial of the time-unrestricted version of the thesis) collapses into radical anti-intentionalism, Intrapersonal Time-Unrestricted Intentionalism collapses into Interpersonal Intentionalism. Moderate anti-intentionalism is thus an inherently unstable position; the two stable positions on this topic are radical anti-intentionalism and intentionalism.²⁷ It is hard to be a moderate anti-intentionalist—even a quietist one.

Though intentionalism has received several powerful defenses in recent years, the battle over counterexamples has been fought, almost exclusively, on the territory of the intentionalists: the question has been whether or not there is a version of intentionalism which can make sense of problematic sorts of perceptual experiences or challenging thought-experiments.²⁸ The present argument is an attempt to turn the tables, and argue that there is no version of anti-intentionalism that does not land in one or another absurdity, by entailing the possibility of one of or more of the scenarios described above.

²⁶ This sort of moderate anti-intentionalist must claim that there are properties relevant to the determination of the contents of the states of a subject which, by their nature, cannot change over the course of that subject's life. The obvious thought here is that something to do with the subject's evolutionary history is relevant. No matter how long a subject lives, one might think, the purposes for which his states evolved cannot change. I'm skeptical that this way out can work. For one thing, the plausibility of evolutionary theories of content seem to decrease when we consider sufficiently long-lived and protean organisms. Further, it's hard to see how these sorts of theories of content can avoid the problem discussed above: that two distinct types of phenomenal states could have evolved to represent the same color property, which would make cases of psychedelic phenomenology + constant content possible. However, I don't think that anything I've said shows definitively that no version of this response on the part of the moderate anti-intentionalist who wants to endorse Intrapersonal Time-Unrestricted Intentionalism will work.

²⁷ This may make some want to rethink the earlier claim that, for example, cases of psychedelic phenomenology + constant content are impossible. Maybe the right thing for an anti-intentionalist to say is just that, contrary to our initial intuitions, these cases really are possible. I've said nothing against this position other than the statement of intuitions about cases above.

²⁸ For defenses of intentionalism, see especially, Harman (1990), Byrne (2001), and Tye (2002). For an illuminating discussion of types of perceptual experience which have been thought to be problematic for intentionalism, see Chap. 4 of Tye (2000).

One might think that the present argument is incomplete without some response to the Lockean argument against Interpersonal Intentionalism above. There are different versions of the Lockean argument—which include different arguments for the conclusion that neither of the inverted subjects is misrepresenting the colors of things. Consideration of all of these is well beyond the scope of this paper. But there is something to be said about the 'parity' argument above: no one who thinks that radical skeptical scenarios are possible should be convinced by it. In a skeptical scenario of almost any sort, the subject would arguably be reasonable to think that the skeptical scenario does not obtain, for just the same reasons that we think that we are reasonable in thinking that we aren't living in some such scenario. But we shouldn't infer from this point that there's something impossible about radical skeptical scenarios. Just so, we shouldn't infer from the correct point that a spectrum-inverted subject would have good reason to believe that she was not misrepresenting the colors the conclusion that it is impossible that she is.

5 Consequences

Two restrictions on this argument should be noted. First, as noted above, it shows that spectrum inversion without difference in representation is impossible, rather than that inversion without misrepresentation is impossible. If it should turn out that spectrum inverted subjects represent the color of the marigold differently, but that, due to the metaphysics of color, neither *mis*represents it, then the possibility of Lockean inversion without misrepresentation is consistent with the foregoing. The argument concerns the relationship between phenomenal character and representational content rather than the metaphysics of what is represented.

Second, the above argument, if successful, shows that *color* phenomenology supervenes on *color* content, not that total visual phenomenology supervenes on total visual content, and still less that total perceptual phenomenology supervenes on total perceptual content. It is a reasonable, but nontrivial, assumption that if intentionalism holds for the case of representation of color, it should also hold for the perceptual representation of other sorts of properties. Indeed, it seems that there is a natural extension of the present form of argument to other cases. Consider, for example, tactile representation of surfaces as rough or smooth, and the phenomenal characters of the tactile experiences which, for us, represent those properties. We can raise the same questions as above about the relations between these experience types. As above, it seems incoherent that we could have a case in which a subject's tactile phenomenology switched rapidly back and forth from smooth-feeling to rough-feeling, while throughout the subject's tactile experience represented the surface as perfectly smooth; so, as above, the combination of psychedelic content with constant content is impossible. From this point, the argument against various forms of anti-intentionalism about tactile representation of texture is parallel to the argument against anti-intentionalism about visual representation of color. I am inclined to think that we could provide parallel arguments for any case of perceptual representation of properties which has an associated phenomenology, and so that this form of argument generalizes to an argument for intentionalism about all modalities of perceptual experience; but that is not a conclusion for which I've directly argued here.²⁹

There are also some surprising extensions of the present sort of argument. The above argument is couched as an argument against anti-intentionalism; but it also, if successful, rules out some forms of intentionalism. Consider, for example, the response to examples of spectrum inversion which is due to Sydney Shoemaker. Roughly, this says that (in at least some cases) spectrum inverted subjects will represent the color of the relevant objects as the same, but differ in their visual representation of some other class of properties, sometimes called 'phenomenal

²⁹ This does not entail that every sort of perceptual representation of properties has an associated phenomenology. I think it is plausible that, for example, perceptual experiences can represent objects as belonging to certain natural kinds, even though two experiences can differ in their representation of natural kind properties without there being any difference between the phenomenal characters of the two experiences. For argument, see Speaks (2009), which is based on some examples from Johnston (2004).

properties' or 'appearance properties.'³⁰ Since this sort of view finds a difference in content to correspond to the difference in phenomenal character between spectrum inverted subjects, this take on cases of spectrum inversion is a version of intentionalism. Nonetheless, the foregoing argument is a challenge to this view.

Proponents of this version of intentionalism face exactly the same challenges as anti-intentionalists: they must specify an external relation between phenomenal character and representation of color which allows spectrum inverted subjects (of certain kinds) to agree in their perceptual representation of color but does not permit the possibility of psychedelic phenomenology + constant content, constant phenomenology + psychedelic content, or full phenomenology + no content.

Proponents of appearance properties might also, like the quietist moderate anti-intentionalist, refuse to specify such a relation; but then they face just the same problems as the moderate anti-intentionalist. They have to endorse one or the other of the following two claims:

- Two subjects can have experiences which are different with respect to their representation of appearance properties but alike with respect to representation of color properties, but no one subject can have two experiences different with respect to representation of appearance properties but alike with respect to representation of color properties. (This is the analogue of Time-Unrestricted Intrapersonal Intentionalism.)
- One subject can have experiences which are different with respect to their representation of appearance properties but alike with respect to representation of color properties, so long as these experiences are separated by some interval of time t ; but no one subject can have two experiences separated by a time less than t which are different with respect to representation of appearance properties but alike with respect to representation of color properties. (This is the analogue of Time-Restricted Intrapersonal Intentionalism.)

The former risks collapse into the view that no pair of experiences of two subjects can differ with respect to representation of appearance properties without also differing with respect to representation of color properties, which undermines the purpose of introducing appearance properties in the first place. The latter risks collapse into the analogue of radical anti-intentionalism: the view that *consecutive* experiences of a subject might differ with respect to representation of appearance properties while being alike with respect to representation of color properties. But this is just to grant the possibility of psychedelic phenomenology + constant content.

The proponent of appearance properties might try make plausible the claim that the cases of psychedelic phenomenology + constant content described above are possible on the grounds that these are cases in which psychedelic *appearance property* phenomenology (rather than psychedelic color phenomenology) accompanies constant representation of color. In my view, this redescription does nothing to make plausible the claim that cases of the sort described at the outset are possible, though this is likely something about which opinions will vary. One way to press the

³⁰ For different versions of this view, see Shoemaker (1994, 2000) and Egan (2006).

worry against the proponent of appearance properties is by considering again the inverse scenario, in which an experience has the constant phenomenology which we experience when looking at a charcoal grey surface under constant illumination, but we visually represent the color of the surface as changing. Is it plausible to defend the possibility of cases of this sort by claiming that, in such a case, our wildly varying visual representation of the color of the surface is obscured by our constant representation of appearance properties?

Perhaps the most far-reaching consequence of this form of argument, however, has to do with the possibility of spectrum inversion. As noted above, there is not one question about the possibility of spectrum inversion; rather, there are a series of questions about whether spectrum inversion is possible while a certain class of facts—such as facts about a subject’s behavioral dispositions, functional architecture, or physical constitution—is held fixed. The above argument is an attempt to show that spectrum inversion while holding fixed perceptual *content* is impossible; but that by itself doesn’t entail the impossibility of the various other sorts of spectrum inversion scenarios which have been discussed in the literature.³¹

One thought is that the foregoing makes these debates less interesting than they would otherwise be. For, one might think, the argument above shows that whether or not spectrum inversion is possible, spectrum inversion without a difference in representation isn’t; and it’s only the latter that can be used to pose a problem for intentionalism.

But, if anything, the present line of argument makes debates about the possibility of spectrum inversion more, not less, important. For consider: if spectrum inversion is possible between subjects alike with respect to a class *A* of properties, then we know that, since spectrum inversion without difference in representation is impossible, two subjects can be alike with respect to their *A*-properties even though the contents of their perceptual states differ. This would be enough to show that (in one important sense of ‘supervenes’) representational properties do not supervene on the *A*-properties.

And it’s hard to see how this result could be limited to *perceptual* representation. Suppose that you and I are spectrum inverted relative to each other; then we will differ with respect to the color our visual experiences attribute to a marigold. Suppose that we both take our experiences at face value, forming beliefs which attribute to the marigold the same color property our experience represents the marigold as having. Then it seems that our beliefs about the color of the marigold will differ in content as well.³² Differences at the level of perceptual representation ramify; only in an odd and coincidental sort of case can two subjects differ in their

³¹ For relevant discussions of the possibility of various kinds of spectrum inversion, see Shoemaker (1975, 1981), Tye (1995), Block (1999), Hoffman (2006), Byrne and Hilbert (2006), and Broackes (2007).

³² I’m skirting questions about ‘nonconceptual content’, which I think are ultimately beside the point here. While it is true that the easiest way of reading the present argument involves attributing the same kind of content to perceptual states and to beliefs—an assumption that most proponents of nonconceptual content will want to reject—the only assumption which is strictly required is that in the default case, a difference in the content of a perceptual experience will issue in a difference in the content of the perceptual belief formed by taking that experience at face value.

perceptual representation of the world while being alike with respect to every proposition they entertain, judge, and believe.³³

This provides us with a general tool for arguing against various theories of mental representation. Suppose that a theory explains the content of mental representations in terms of the *A*-properties. Then we can ask: is spectrum inversion possible between subjects alike in their *A*-properties? If so, then mental representation fails to supervene on the *A*-properties, and the theory of mental representation in question is false.

The more specific the theory of mental representation, the more powerful this line of argument. It is notoriously difficult to decide whether spectrum inversion between full functional duplicates is genuinely possible. But it might be easier—to use Jerry Fodor’s well-known and admirably specific theory of mental representation as an example—to decide whether spectrum inversion is possible between subjects who are alike with respect to, for example, the dependence relations among nomological connections between representations in their language of thought and properties in the world.³⁴ If this is possible, then, if the preceding argument is sound, Fodor’s theory of mental representation must be false.

In fact, this sort of argument gets even easier if we note that the argument does not even require full spectrum *inversion*; all that’s needed is spectrum *shift*. For consider the following scenario:

A subject is looking intently at a well-lit surface. Over the course of a few seconds, his experience goes from being (as we would put it) bright-green-feeling to being bright-greenish-yellow feeling to being bright-yellowish-green feeling, and constantly repeats this pattern. But, the whole time, he is visually representing the wall as pure green; it visually seems to him throughout that the wall is pure green.

³³ There’s a connection here with a powerful version of the inverted spectrum argument against intentionalism which can be taken from the discussion in Thau (2002):

Suppose that two subjects spectrum inverted relative to each other, Invert and Nonvert, are members of a single linguistic community. Since they both use, for example, ‘blue’ and ‘yellow’ to apply to the same things, and seem to understand each other perfectly well, we can take it that each uses these words with the same meaning. But they use these words to report their perceptual beliefs; and surely they aren’t mis-reporting the contents of their own beliefs! So, since each says on the basis of their visual experience ‘I believe that that marigold is yellow’, we can safely assume that they have the same beliefs about the colors of things. But they form these beliefs on the basis of their visual experiences of the colors of things; and surely they aren’t mistaken about how their own visual experiences represent the world as being! So we can take it that their visual experiences agree in their representation of the colors of things. But then two subjects can have experiences alike with respect to the representation of color but different with respect to color phenomenology, and Interpersonal Intentionalism (or at least the kind defended above) must be false, after all.

In my view, this is the most challenging version of the inverted spectrum argument against intentionalism. I’m inclined to think that the right response is to reject the supposition that the two use the words ‘blue’ and ‘yellow’ to stand for the same property. Why not think that in some cases two people can use words to stand for different properties, even if this difference in meaning could never come to light?

³⁴ See Fodor (1990).

This is, obviously, a variant on the examples of psychedelic phenomenology + constant content discussed above. It is, perhaps, not as spectacularly impossible as those scenarios; but, I take it, it is fairly clear that this scenario is, still, impossible, and for much the same reasons as the cases discussed above. If this is right, then by argument parallel to the above it follows that spectrum shift without a difference in representation is impossible.³⁵ But, if so, it follows that if two subjects can be spectrum shifted with respect to each other while alike with respect to their A-properties, then mental representation cannot be explained in terms of the A-properties.³⁶

This is bad news for theories of mental representation. A common—perhaps the dominant—attitude in contemporary philosophy of mind has it that while consciousness may remain something of a mystery, we can more or less ignore this mystery while constructing theories of mental representation. Sometimes, this attitude is accompanied by the acknowledgement that facts about the phenomenal characters of a subject's experiences might not even supervene on the sorts of properties adduced to explain mental representation.³⁷ If the above is correct, this combination of views is unstable. If the phenomenal character of perceptual experience does not supervene on the physical facts, then there is no physicalist reduction of mental representation; if there is a physicalist reduction of mental representation, the examples which purport to show the failure of the supervenience of phenomenal character on the physical world must be impossible scenarios. There is no comfortable marriage of confident materialism about mental representation with quiet agnosticism about the metaphysics of the phenomenal character of experience.

Acknowledgments Thanks for helpful comments on previous versions of this paper to Casey O'Callaghan, Adam Pautz, participants in a colloquium at McGill University, and the students in my graduate seminar at Notre Dame in the Spring of 2007.

References

- Alston, W. (2005). Perception and representation. *Philosophy and Phenomenological Research*, 70, 253–289.
- Block, N. (1990). Inverted Earth. *Philosophical Perspectives*, 4, 52–79.
- Block, N. (1995). Mental paint and mental latex. *Philosophical Issues*, 7, 19–49.
- Block, N. (1999). Sexism, racism, ageism and the nature of consciousness. *Philosophical Topics*, 26(1–2), 39–70.
- Block, N. (2003). Mental paint. In M. M. Hahn & B. Ramberg (Eds.), *Reflections and replies: Essays on the philosophy of Tyler Burge*. Cambridge, MA: MIT Press.

³⁵ This is a kind of reversal of the powerful argument against intentionalism in Block (1999), and is connected to recent discussion of the puzzle of 'true blue' in, for example, Tye (2006). The response to these arguments suggested by the foregoing is pretty much the one in Byrne and Hilbert (2007); see also Tye (2007).

³⁶ Since there are actual cases of spectrum shifted subjects, this gives us a test for proposed theories of mental representation. To again use Fodor's theory as an example, do every pair of spectrum shifted subjects differ with respect to, for example, dependence relations among nomological connections between mental representations and instantiations of color properties?

³⁷ See, for example, Kim (2007) and Chalmers (1996).

- Boghossian, P., & Velleman, D. (1989). Colour as a secondary quality. *Mind*, 98(389), 81–103.
- Brewer, B. (2006). Perception and content. *European Journal of Philosophy*, 14(2), 165–181.
- Brewer, B. (2008). How to account for illusion. In A. Haddock & F. Macpherson (Eds.), *Disjunctivism: Perception, action, knowledge* (pp. 168–180). New York: Oxford University Press.
- Broackes, J. (2007). Black and white and the inverted spectrum. *Philosophical Quarterly*, 57(227), 161–175.
- Byrne, A. (2001). Intentionalism defended. *Philosophical Review*, 110(2), 199–240.
- Byrne, A., & Hilbert, D. (Eds.). (1997). *Readings on color: Volume 1, The philosophy of color*. Cambridge, MA: MIT Press.
- Byrne, A., & Hilbert, D. (2006). Hoffman's "proof" of the possibility of spectrum inversion. *Consciousness and Cognition*, 15, 48–50.
- Byrne, A., & Hilbert, D. (2007). Truest blue. *Analysis*, 67, 87–92.
- Campbell, J. (2002). *Reference and consciousness*. New York: Oxford University Press.
- Chalmers, D. (1996). *The conscious mind: In search of a fundamental theory*. New York: Oxford University Press.
- Chisholm, R. (1957). *Perceiving: A philosophical study*. Ithaca, NY: Cornell University Press.
- Cohen, J. (2003). Color: A functionalist proposal. *Philosophical Studies*, 113(1), 1–42.
- Cohen, J. (2004). Color properties and color ascriptions: A relationalist manifesto. *The Philosophical Review*, 113(4), 451–506.
- Cohen, J. (2007). A relationalist's guide to error about color perception. *Noûs*, 41(2), 335–353.
- Cohen, J. (2008). Colour constancy as counterfactual. *Australasian Journal of Philosophy*, 86(1), 61–92.
- Egan, A. (2006). Appearance properties? *Noûs*, 40(3), 495–521.
- Fodor, J. (1990). A theory of content, I: The problem. In *A theory of content and other essays* (pp. 51–88). Cambridge, MA: MIT Press.
- Harman, G. (1987). (Nonsolipsistic) conceptual role semantics. In G. Harman (Ed.), *Reasoning, meaning, and mind* (pp. 206–232). Oxford: Clarendon Press.
- Harman, G. (1990). The intrinsic quality of experience. *Philosophical Perspectives*, 4, 31–52.
- Hilbert, D. (2005). Color constancy and the complexity of color. *Philosophical Topics*, 33(1), 141–158.
- Hilbert, D., & Kalderon, M. (2000). Color and the inverted spectrum. In S. Davis (Ed.), *Color perception: Philosophical, psychological, artistic, and computational perspectives*. New York: Oxford University Press.
- Hoffman, D. (2006). The scrambling theorem: A simple proof of the logical possibility of spectrum inversion. *Consciousness and Cognition*, 15, 31–45.
- Johnston, M. (1992). How to speak of the colors. *Philosophical Studies*, 68(3), 221–263.
- Johnston, M. (2004). The obscure object of hallucination. *Philosophical Studies*, 120, 113–183.
- Kalderon, M. E. (2007). Color pluralism. *Philosophical Review*, 116(4), 563–601.
- Kim, J. (2007). *Physicalism, or something near enough*. Princeton, NJ: Princeton University Press.
- Macpherson, F. (2005). Colour inversion problems for representationalism. *Philosophy and Phenomenological Research*, 70(1), 127–151.
- Marcus, E. (2006). Intentionalism and the imaginability of the inverted spectrum. *Philosophical Quarterly*, 55(224), 321–339.
- McLaughlin, B. (2003). Color, consciousness, and color consciousness. In Q. Smith (Ed.), *New essays on consciousness* (pp. 97–154). Oxford: Oxford University Press.
- Peacocke, C. (1983). *Sense and content: Experience, thought, and their relations*. Oxford: Oxford University Press.
- Shoemaker, S. (1975). Functionalism and qualia. *Philosophical Studies*, 27, 291–315.
- Shoemaker, S. (1981). The inverted spectrum. *Journal of Philosophy*, 74(7), 357–381.
- Shoemaker, S. (1994). Phenomenal character. *Noûs*, 28, 21–38.
- Shoemaker, S. (2000). Phenomenal character revisited. *Philosophy and Phenomenological Research*, 60, 465–467.
- Shoemaker, S. (2001). Introspection and phenomenal character. *Philosophical Issues*, 28(2), 247–273.
- Speaks, J. (2009). Transparency, intentionalism, and the nature of perceptual content. *Philosophy and Phenomenological Research*, 79(3), 539–573.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.
- Stampe, D. (1979). Toward a causal theory of linguistic representation. In P. French, T. Uehling, & H. Wettstein (Eds.), *Contemporary perspectives in the philosophy of language* (pp. 81–102). Minneapolis: University of Minnesota Press.
- Thau, M. (2002). *Consciousness and cognition*. New York: Oxford University Press.

- Travis, C. (2004). The silence of the senses. *Mind*, 113, 57–94.
- Tye, M. (1995). *Ten problems of consciousness: A representational theory of the phenomenal mind*. Cambridge, MA: MIT Press.
- Tye, M. (2000). *Consciousness, color, and content*. Cambridge, MA: MIT Press.
- Tye, M. (2002). Representationalism and the transparency of experience. *Nous*, 36, 137–151.
- Tye, M. (2006). The puzzle of true blue. *Analysis*, 66, 173–178.
- Tye, M. (2007). True blue redux. *Analysis*, 67, 92–93.