

DON'T PANIC: Tye's intentionalist theory of consciousness*

Alex Byrne, MIT

Consciousness, Color, and Content is a significant contribution to our understanding of consciousness, among other things. I have learned a lot from it, as well as Tye's other writings. What's more, I actually *agree* with much of it—fortunately for this symposium, not all of it.

The book continues the defense of the “PANIC” theory of phenomenal consciousness that Tye began in *Ten Problems of Consciousness* (1995). A fair chunk of it, though, is largely independent of this theory: the discussion of the knowledge argument, the explanatory gap, and color. Tye says much of interest about these topics. But as most of my disagreement is with the PANIC theory, I shall concentrate on that.

The PANIC theory is nothing short of ambitious. It is a reductive account of phenomenal consciousness in intentional/functional terms. Tye further gives, at least in outline, a broadly physicalistic account of intentionality (a “naturalized semantics”), in terms of causal covariation. Putting the PANIC theory and Tye's naturalized semantics together, the result is a physicalistically acceptable theory of phenomenal consciousness.

The two parts of this package are independent. A naturalized semantics can be combined with dualism about consciousness (a position close to this is in Chalmers 1996). And a PANIC theorist is at liberty to endorse a rival physicalistic theory of intentionality, or indeed could take intentionality to be entirely irreducible.

* Many thanks to Michael Glanzberg, Ned Hall, Sally Haslanger, Jim John, Sarah McGrath, Jim Pryor, Mark Richard, Susanna Siegel, Robert Stalnaker, Ralph Wedgwood, and Steve Yablo.

The plan is this. Section 1 briefly airs a concern about Tye's naturalized semantics. The rest of the paper focuses on the PANIC theory. One important component of Tye's view, discussed in section 2, is *intentionalism*—roughly, the claim that the phenomenal character of an experience is fixed by its propositional content. Intentionalism is controversial enough, but the PANIC theory (explained in section 3) is considerably stronger. The various additions the PANIC theory makes to intentionalism are discussed in sections 4, 5, and 6. Finally, section 7 sketches a couple of alternative suggestions for treating some of the problems raised in the preceding three sections.

Before getting down to business, some terminology needs to be clarified.

The *phenomenal character* of an experience can be introduced by examples: the experience of tasting sugar differs in phenomenal character from the experience of tasting lemon juice; the experience of seeing ripe tomatoes differs in phenomenal character from the experience of seeing unripe ones; your experience and the corresponding experience of your twin on Twin Earth have the same phenomenal character; if Invert is “spectrally inverted” with respect to Nonvert, then Invert's tomato-experiences differ in phenomenal character from Nonvert's; and so on. Note that on the usage adopted here, the phenomenal character of an experience is a *property of* the experience; sometimes ‘qualia’ is used equivalently, but sometimes not (see, for example, Lycan 1996, 69-70).

The *propositional content*—or, simply, content—of an experience captures the way the world *perceptually seems* to the subject of the experience. When one looks at a purple pentagon in good light, it seems *that there is a purple pentagon before one*. Clearly the proposition that there is a purple pentagon before one falls short of completely characterizing the way the world seems, but pretend otherwise for illustration. If there *isn't* a purple pentagon before one, then the content of the experience is *false*, and the experience is some kind of *illusion*. If there *is* a purple pentagon before one, then the content of the experience is *true*, and the experience is *veridical*.

The content of experience, or perceptual content, can be also introduced in a more familiar idiom. Perceptual experiences are species of *propositional attitude*: it visually (aurally/tactually, etc.) appears that p . If it visually appears that p (and if the proposition that p completely characterizes the way things visually appear), then the content of one's experience is just the proposition that p .¹

There are many hard questions concerning perceptual content. Imagine someone with normal vision looking at an object that is shaped and colored exactly like a yellow lemon. She might describe the scene by saying that there seems to be a yellow ripe lemon before her. Presumably the content of her experience at least concerns the color and shape of the object. But does it also specify the object before her as *ripe*, or as a *lemon*? Is her experience some kind of *illusion* if the object is a yellow but unripe lemon, or if the object is made of papier-mâché? Would the content of her experience be different if a qualitatively identical but numerically distinct object were before her eyes? Connectedly, would the content of her experience be the same, or at least importantly similar, if she were hallucinating a lemon?

Evidently the notions just introduced—the *phenomenal character* and *content* of an experience—are not especially clear; however, I assume with Tye that they are clear enough to support some theorizing.

¹ More exactly: if it visually appears that p at time t (and if the proposition that p completely characterizes the way things visually appear), then the content of one's experience at t is the proposition that p . This complication will be ignored. Note that 'It visually appears that p ' is a piece of semi-technical terminology. Whether the proposition that Tye is friendly could be the content of one's visual experience is not to be settled by considering the use of the English sentence 'It visually appears that Tye is friendly'.

Finally, a cautionary-cum-apologetic note. Partly to make the discussion fit smoothly with various quoted passages, *events* (for instance, experiences, and episodes of thinking), and *states* (for instance, beliefs), will be lumped together as *states*.²

1. Tye's naturalized semantics

Tye's causal covariational account of intentionality is this:

[Sensory state] *S* represents that *P* =df If optimal conditions were to obtain, *S* would be tokened in [creature] *c* if and only if *P* were the case; moreover, in these circumstances, *S* would be tokened in *c* because *P* is the case. (2000, 136, note omitted; cf. 1995, 101)³

“Optimal conditions” are explained as follows:

In the case of evolved creatures, it is natural to hold that such conditions for vision involve the various components of the visual system operating as they were designed to do in the sort of external environment in which they were designed to operate. (138)

It seems to me that Tye himself has supplied compelling counterexamples against this proposal, namely various perceptual illusions, in particular the Müller-Lyer illusion

² I am not pretending that this policy is entirely harmless. For a useful critical discussion of “states” and other ontological categories in the philosophy of mind, see Steward 1997.

³ As he says (1995, 101) this account derives from Stampe 1977 and Stalnaker 1984. Tye later adds a complication (2000, 139-40) in the style of Fodor's asymmetric dependency account (Fodor 1990, ch. 4); this is not relevant here.

(1995, 102; 2000, 106). In the latter illusion, one's visual experience represents (incorrectly) that the lines are of different lengths, even in conditions that are presumably optimal.

It might be replied that the two-dimensional Müller-Lyer diagram is not supposed to be included in the “sort of external environment” in which the components of the visual system were “designed to operate”. If so, we need much more of a story about the right kind of external environment than Tye supplies. And in any case, this reply does not work: illusions like the Müller-Lyer occur when viewing ordinary three-dimensional scenes (DeLucia and Hochberg 1991). If “optimal conditions” are to play a central role in a naturalized semantics, they need to be explained along quite different lines.⁴

2. Intentionalism

Setting Tye's naturalized semantics aside, let us begin our investigation of the PANIC theory. According to Tye, “necessarily, experiences that are alike in their representational contents are alike in their phenomenal character” (2000, 45), a thesis he calls *representationalism*. The PANIC theory is supposed to be a version of representationalism (2000, x, 45). If representationalism is correct, the phenomenal difference between experiences in *different* sensory modalities—between seeing and hearing, for example—is due to a difference in content. But one might be more cautious. Tye distinguishes representationalism from a “modality-specific, weak representational thesis R”:

⁴ Essentially the same problem arises for Dretske's (1995) theory of naturalized semantics (which leans more heavily than Tye's on teleology: see Tye 2000, 119). Dretske discusses this problem in an endnote (174, n. 13), and gives a version of the reply mentioned above.

Necessarily, visual experiences that are alike with respect to their representational contents are alike phenomenally. (2000, 69)

For present purposes the PANIC theory needs to be sharply separated from both representationalism and “thesis R”. To avoid confusion it is best to introduce some different terminology.

Intramodal intentionalism is the claim that, *within* a perceptual modality, the phenomenal character of an experience supervenes on its content. An intramodal intentionalist therefore holds thesis R and its analogue for the other senses (which may be taken to include uncontroversial examples like olfaction and audition). *Intermodal intentionalism* is the claim quoted at the start of this section: necessarily, experiences alike in representational content are alike in phenomenal character. Hence, *intermodal intentionalism* implies *intramodal intentionalism*, but not conversely. These two sorts of intentionalism are *unrestricted* just in case they encompass not just paradigmatic perceptual experiences, but also sensations, like pain and nausea.⁵

To illustrate the core of these intentionalist positions, imagine that Invert is “spectrally inverted” with respect to Nonvert. They are both looking at a tomato, and because of the inversion their experiences differ in phenomenal character. Despite this difference, might Invert’s and Nonvert’s experiences have exactly the same content (they both represent the tomato as red, etc.)? According to some philosophers—notably Block (1990, 2000)—the answer is yes, while intentionalists disagree.⁶ Again, some

⁵ Lycan (1996) is an example of an *intramodal* intentionalist (according to him, functional role, not content, accounts for the phenomenal difference between sensory modalities); McGinn (1991, ch. 2) is an example of a *restricted* intentionalist (he thinks sensations have no content). This terminology is taken from Byrne 2001.

⁶ Other anti-intentionalists include Burge (forthcoming), Levine (2001), and Peacocke (1983).

philosophers argue that a “zombie” is possible: a creature intentionally identical to you or me, but whose “experiences” have no phenomenal character: it visually appears to her, say, that there is a pink circle ahead, but there is nothing it’s like for her to enjoy this experience.⁷ Intentionalists deny that any such zombie is possible.

Intentionalism is obviously controversial, and Tye’s brand—intermodal unrestricted intentionalism—is even more so. As it happens, I agree with Tye that intermodal unrestricted intentionalism is correct (Byrne 2001); ‘intentionalism’ will henceforth be used for this strong thesis, unless the context indicates otherwise.

Now some mental states have content, but do not have phenomenal character. For example, there is nothing it’s like to believe that today is Wednesday—or, at any rate, there need be nothing it’s like to have this belief (one may have it during one’s lunchtime nap). More controversially, there need be nothing it’s like to recall (consciously) that today is Wednesday, or to wonder (consciously) whether today is Wednesday. At any rate, wondering whether today is Wednesday is hardly, to borrow a phrase of Block’s, “phenomenologically impressive”.

So a question naturally arises: what is the difference between those intentional states that have phenomenal character and those that don’t? What is the ingredient X that makes an intentional state one with phenomenal character? This is a question for both the

⁷ In the usage of this paper, when a subject undergoes an “experience” with the content that *p*, it perceptually appears *to her* that *p*. If some sub-personal state of the subject has the content that *p* (and so it does not appear *to her* that *p*), then this state is not an experience. Therefore the perceptual states of certain blindsight patients are not experiences (it does not appear *to the subject* that there is an ‘O’ before her). Note that this usage does *not* trivialize the claim that all experiences have phenomenal character. The zombie possibility mentioned above is supposed to be a case where the subject has an experience (in the sense used here), but with no phenomenal character.

I think this usage of ‘experience’ is (in this respect) pretty close to Tye’s, but it is certainly not universal in the literature. For a broader use of ‘experience’ that includes blindsight cases, see Carruthers 2000, ch. 6.

intentionalist and his opponent. An anti-intentionalist may say something entirely unhelpful (like “Qualia”), or he may offer something more substantive, for instance a theory of “sensational properties” (Peacocke 1983). It is important to emphasize that the intentionalist is not under any greater obligation: a substantive reply is desirable, but not mandatory.

Comparing intentionalism with other supervenience theses helps to reinforce the point. Take, for example, the claim that the mental supervenes on the physical (say, a global supervenience thesis of the sort in Braddon-Mitchell and Jackson 1996, ch. 1). Given that this world contains minds, the supervenience thesis tells us that any physical duplicate of this world also contains minds. Consistently with this, it might be quite obscure why this world contains minds at all. Why does *this* arrangement of atoms in the void necessitate the existence of minds? What is the ingredient X that turns mere matter into *thinking* matter? Supervenience theses do not give satisfying answers to such questions. For another example, take the claim that the evaluative supervenes on the descriptive. Given that there are evaluative claims true at this world: Jones is brave; Alice ought to give Bert his banana back, etc., the supervenience thesis tells us that these claims are true at any descriptive duplicate of this world. Consistently with this, it might be quite obscure why *these* descriptive claims necessitate such-and-such evaluative claims.

Lovers of mystery have nothing to fear, then, from supervenience; in particular, those who find consciousness especially perplexing need not spurn intentionalism.⁸

The PANIC theory, as we will see in the following section, goes considerably beyond intentionalism: it supplements it with a substantive proposal for the philosopher’s stone, the elusive ingredient X.⁹

⁸ As McGinn (1991) clearly recognizes.

3. *The PANIC theory*

The PANIC theory is this: “phenomenal character is one and the same as Poised, Abstract, Nonconceptual, Intentional Content” (2000, 63; cf. 1995, 137).

Three bits of PANIC terminology need to be explained: ‘poised’, ‘abstract’, and ‘nonconceptual’ (“intentional content” is just propositional content, a.k.a. representational content). Take ‘abstract’ first. This applies in the first instance to propositions or contents. A proposition is abstract iff it is not object-dependent (1995, 138; 2000, 62). Thus the proposition *that Tye is a philosopher* is not abstract, because its truth at any circumstance of evaluation depends on how things are with a particular individual, viz. Tye. The propositions *that (some x) x is a philosopher* and *that (the x: x is a man drinking a martini) x is a philosopher*, on the other hand, are abstract. We can speak derivatively of an abstract mental state: a *state* is abstract iff its content is abstract. For example, the belief that (some x) x is a philosopher is abstract.

Now turn to ‘poised’. This applies in the first instance to mental states, not to contents. A state is poised iff it “stand[s] ready and available to make a direct impact on beliefs and/or desires” (2000, 62; cf. 1995, 138). A visual experience as of a tomato is poised, because it typically causes a belief about the tomato “if attention is properly focused” (62). However, earlier stages of visual processing that represent, say, “changes in light intensity” are not poised: “the information they carry is not directly accessible to the relevant cognitive centers” (2000, 62). We can speak derivatively of poised contents: a content is poised iff it is the content of some poised state.

Finally, ‘nonconceptual’. This is the most problematic of the three, and many pixels will be spilt on it later (section 6). But for now, we can make do with the following

⁹ Naturalistic theories of consciousness in this style—intentionalism + X—are very popular. See, in particular, Carruthers 2000, Dretske 1995, Kirk 1994, Lycan 1996.

explanation: “The claim that the contents relevant to phenomenal character must be *nonconceptual* is to be understood as saying that the general features entering into these contents need not be ones for which their subjects possess matching concepts” (1995, 139). A *state* is nonconceptual iff it has nonconceptual content.

So much for PANIC, but now something needs to be said about phenomenal character. Tye intends the equation ‘Phenomenal character is PANIC’ to be understood as *identifying* phenomenal character with a certain kind of content: “phenomenal character is one and the same as representational content that meets certain further conditions” (2000, 45). As I understand it, this “representational content that meets certain further conditions” is the *content of experience*, as explained at the start of this paper.¹⁰ On Tye’s usage, then, the phenomenal character of my visual experience just *is* the content of my

¹⁰ Since misunderstanding might set in at exactly this point, some extra clarification can’t hurt. Consider the following passage from Tye:

The term ‘experience’ can be used in broader and narrower ways. I have assumed in my remarks above that it is correct to say that we have visual experiences as of coins, telescopes, and so forth. Some may prefer to restrict the term ‘experience’ to states with nonconceptual content, counting the rest as judgments superimposed upon experience proper. The issue seems to me purely terminological. I am here adopting the broader usage... (2000, 76; cf. 1995, 140 on “experiential episodes, broadly construed”).

This paper adopts Tye’s *narrow* use of ‘experience’, or near enough. The content of experience, on the narrow use of ‘experience’, goes hand in hand with the intuitive conception of a perceptual *illusion*: the content of an experience is false iff the experience is an illusion (perhaps just a partial one). So, on this narrow use, it’s clear that we sometimes have visual experiences that represent objects *as purple*; it’s false, or at least controversial, that we have visual experiences that represent objects *as poisonous* (cf. 2000, 54-5); and it’s uncontroversially false that we have visual experiences that represent objects *as friends of Tye*. It should be emphasized that ‘experience’, as used here, is not *defined* to apply only to states with nonconceptual content. Whether experience has nonconceptual content is a substantive issue.

experience: a particular *content* or *proposition* that is also abstract, poised, and nonconceptual.

I myself find this usage a bit confusing. On the way Tye sets things up in chapter 3 of *Consciousness, Color, and Content*, the investigation of the relation between phenomenal character and content begins before we have even settled whether the phenomenal character of an experience is a *property*. The hypothesis that “visual phenomenal character” is a quality (i.e. property), specifically a “quality of the surface experienced”, is considered and rejected (48). The conclusion of the investigation is that (visual) phenomenal character is *not* a property; rather it is a kind of content.

It seems to me preferable to sort out these basic ontological questions first, before starting the philosophical argument. And this is best done, I think, by stipulating that the phenomenal character of an experience *e* is a *property*, specifically a *property of e*: that property that types *e* according to what it’s like to undergo *e*. (This sort of account was given at the start of this paper.) On this alternative and fairly common usage, although the result of an investigation might be that phenomenal characters were, say, functional or physical properties, it *couldn’t* turn out that they were propositions, and so not properties at all.

For these reasons, the PANIC theory will be set out here with phenomenal character understood as a *property of* a mental state, *a fortiori* not a proposition. More specifically, in the usage of this paper, the phenomenal character of a mental state is that *maximally determinate* property that types the state in respect of what it’s like to be in the state. That is, e_1 and e_2 have the same phenomenal character iff what it’s like to undergo e_1 is exactly what it’s like to undergo e_2 . (We should add the stipulation that if there is nothing it’s like to be in *e*, then *e* has *no* phenomenal character.) On this conception, the phenomenal character of the experience of looking at a tomato is *different* from the phenomenal character of the experience of looking at raspberry (despite the fact that they

have something phenomenal in common), and the phenomenal character of your experience is the *same* as that of your twin on Twin Earth.

Tye's identification of phenomenal character with PANIC can now be unpacked as follows. Let S be a mental state with phenomenal character Q . On Tye's view, the intentional content of S will be both abstract and nonconceptual.¹¹ Let it be the proposition P . Then:

$Q =$ the property of being poised, and of having abstract nonconceptual content P .

Let us call this general thesis *PANIC*. It is equivalent to the PANIC theory, assuming I have understood the latter correctly.

Notice that **PANIC** implies that if two states have the *same* phenomenal character, then they have the same content. So, for example, since my visual experience when I see Tye at a conference has the same phenomenal character as my twin's visual experience when he sees twin-Tye on some duplicate of Earth, according to **PANIC** our two experiences have the same content. And it is a very short step from this to the conclusion that perceptual content is not object-dependent; that is, to the conclusion that perceptual content is "abstract". (The content of my experience can hardly involve Tye, because my twin's doesn't, and his experience is supposed to have the *same* content.) In other words, the simpler equation ' $Q =$ the property of being poised, and of having nonconceptual content P ' implies, with minimal further assumptions, the longer one displayed above. The A part of the PANIC theory is therefore not an optional extra.

¹¹ I am here completely ignoring Tye's "broad usage" of 'experience' (see preceding footnote). On that usage, and according to Tye, the content of an experience often won't be (entirely) abstract or nonconceptual.

What is the relation between the PANIC theory (i.e. **PANIC**) and (intermodal, unrestricted) intentionalism? Clearly intentionalism does not imply **PANIC**. An intentionalist may deny, for instance, the following consequence of **PANIC**—that any state with phenomenal character is poised. However, as Tye in effect notes, **PANIC** does imply intentionalism. To see this, let e_1 and e_2 be experiences with, respectively, contents P_1 and P_2 , and characters Q_1 and Q_2 , and assume that **PANIC** is true. Then:

Q_1 =the property of being poised, and of having abstract nonconceptual content P_1 .

And:

Q_2 = the property of being poised, and of having abstract nonconceptual content P_2 .

Therefore, if Q_1 and Q_2 are distinct, so are P_1 and P_2 . Hence, given **PANIC**, intentionalism follows: if any two possible experiences differ in phenomenal character, they differ in content.

According to **PANIC**, an intentional state lacks phenomenal character just in case it isn't poised, or doesn't have abstract or nonconceptual content. So Tye's proposal for ingredient X—the ingredient that makes an intentional state one with phenomenal character—is P + A + N.

If $X = P + A + N$, then the significance of this discovery can hardly be exaggerated. Unfortunately, as is argued in the next three sections, there are major problems with each of P, A, and N.

4. Poisedness

Tye explains the notion of a state's being poised as follows:

This condition is essentially a functional role one. The key idea is that experiences and feelings, qua bearers of phenomenal character, play a certain distinctive functional role. They arise at the interface of the nonconceptual and conceptual domains, and they stand ready and available to make direct impact on beliefs and/or desires. For example, how things phenomenally look typically causes certain cognitive responses—in particular, beliefs as to how they are if attention is properly focused. Feeling hungry likewise has an immediate cognitive effect, namely the desire to eat. In the case of feeling pain, the typical cognitive effect is the desire to protect the body, to move away from what is perceived to be producing pain. And so on. States with nonconceptual content that are not so poised lack phenomenal character. (2000, 62)

On the PANIC theory, an experience that is not poised has no phenomenal character, and this is Tye's explanation of why there's nothing it's like for the blindsight subject to see an 'O'-shaped figure, even though she can reliably identify it as such. In such subjects, "there is no complete, unified representation of the visual field, the content of which is poised to make a direct difference in *beliefs*. Blindsight subjects do not believe their guesses. The cognitive processes at play in these subjects are not belief-forming at all" (2000, 63).

The poisedness requirement is quite weak. As I understand it, a pang of hunger, say, is poised just in case it stands "ready and available to have a direct impact" on *some* beliefs and/or desires—which need *not* include "the desire to eat". And this is just as well, because it is perfectly possible to feel hungry while having no tendency to want to eat (a state dieters strive for). And afterimage experiences do not typically cause beliefs "as to how things are" (that is, beliefs that endorse the content of the experience). When one has a green circular afterimage experience, one does not typically believe that there is a green circular film floating before one. However, if the experience stands "ready and

available” to cause some *other* belief—say, the belief that something is wrong with one’s eyes—then it will be poised. Again, take the “waterfall illusion” (2000, 75). This arguably involves an experience with an inconsistent content, that the rocks by the side of the waterfall are both moving and not moving. The experience does not typically cause the belief that the rocks are both moving and not moving, and yet it is certainly supposed to be poised.

Although poisedness may well be a *necessary* condition for phenomenal character, it does not seem to turn A+N into a *sufficient* condition. Consider the cortically blind patient described by Mestre et al. (1992), who can discriminate “optic flow” (the changes in the retinal array produced by the organism’s motion).¹² He can use his “blindsight” to navigate past obstacles in a cluttered environment, and so something occurs in him that plays part of the information processing role of visual experiences—let us say he has *quasi*-experiences.¹³ We may assume that his quasi-experiences are abstract and non-conceptual. So, on the PANIC theory, their lack of phenomenal character must be traced to the absence of poisedness. Surely, though, the subject’s quasi-experiences *are* poised. They cause the appropriate beliefs: if the subject didn’t have beliefs about various obstacles in his path, he wouldn’t be able to avoid them. Admittedly, the subject cannot, in the normal spontaneous fashion, *verbally express* these beliefs. But that does not mean that he does not have them: one’s beliefs may manifest themselves in one’s

¹² See Milner and Goodale 1995, 85. I am indebted to Carruthers’ (2000, 154-68) discussion of this and other examples; he puts them to a related but somewhat different use.

¹³ See note 7 above.

non-verbal behavior. If beliefs are Ramsey's "maps by which we steer", then the cortically blind patient has the appropriate beliefs about his environment.¹⁴

It might be replied that there are two sorts of beliefs (and desires), and that the poisedness requirement relates to only one kind. First, there are beliefs/desires that are available for use in practical and/or theoretical reasoning, and reportable in speech.¹⁵ Second, there are beliefs/desires that (merely) interact with each other to control bodily movement. And if the "beliefs and/or desires" mentioned in the poisedness requirement are solely of the *first* kind, then the cortically blind patient's quasi-experiences are *not* poised.

It isn't likely that Tye would endorse this reply (cf. 2000, ch. 8, on the beliefs of simple animals). And in any case, it just isn't clear why poisedness defined in terms of the *first* sort of belief/desire is the crucial phenomenology-maker. Given that poisedness defined in terms of the *second* sort of belief/desire fails to turn A+N into a sufficient condition, why should we be so confident that a definition in terms of the first sort does any better? (A similar complaint is nicely developed in Carruthers 2000, ch. 6.)

What's more, poisedness defined in terms of the first sort of belief/desire does *not* seem to turn A+N into a sufficient condition. Remember that the poisedness requirement is apparently quite weak: no constraint is placed on the contents of the beliefs or desires that a poised state stands "ready and available" to cause. Imagine someone rather like a

¹⁴ Some caution is needed. The patient was not completely blind, having a small amount of macular and perifoveal sparing. Mestre et al. report that "motion perception, as evaluated with optical flow patterns, appeared to be functional in perimetrically blind parts of his visual field", and conclude: "These results support the hypothesis that the ability to visually control locomotion was preserved in the blind parts of his visual field. We cannot, however, exclude the possibility of a fundamental contribution of his residual intact visual field to his ambulatory autonomy. Nonetheless, it seems unlikely that his capacities were based only on this residual field" (1992, 791).

¹⁵ That is, beliefs and desires that are (something like) "access-conscious" in the sense of Block 1995.

blindsight patient, who is looking at a tomato, and who is in a state *S* with the content of a normal visual experience as of a ripe tomato. The subject does not have the beliefs (at least of the first sort) that are typically produced by an experience as of a ripe tomato. The subject says he doesn't see anything; he won't reach out if asked to pick up the nearest tomato; and so on. However, due to some quirk of his inner wiring, his state *S* does cause the desire to eat. "I'm famished", he spontaneously says, when facing a ripe tomato, and tucks enthusiastically into the hamburger pressed into his hands. Therefore, if poisedness is defined in terms to the first sort of belief/desire, his state *S* is poised. Moreover, since the content of *S* is the same as that of a normal experience as of a tomato, and since (as noted in section 3 above) the PANIC theory entails intentionalism, it follows that the subject is enjoying a phenomenally conscious experience as of a tomato. That is not credible.

5. *Abstractness*

As explained in section 3, a proposition is abstract iff it is object-independent. According to Tye, when one perceives a certain ripe tomato *o*, for example, the content of one's experience is not an object-*dependent* proposition—say, that *o* is red and round—but instead an object-*independent* proposition—say, that (some *x*) *x* is red and round.¹⁶

Also noted in section 3 was the point that once an equation along the lines of 'Phenomenal character *Q* = the property of having nonconceptual content *P* and...' has been established, then the A part of the PANIC theory comes along (almost) for free. That is, the identity thesis, together with the very plausible assumption that the representation of a *particular individual* (e.g. Tye as opposed to Twin-Tye) makes no

¹⁶ For a qualification about the use of 'experience', see footnotes 10 and 11 above. For the purposes of illustration, this section assumes that perceptual content is linguistically expressible; this will be questioned later in section 7.

distinctive contribution to phenomenal character, implies that perceptual content is abstract.

So, *if* Tye has an argument for an equation of the form ‘Phenomenal character Q = the property of having nonconceptual content P and...’, *without* assuming that perceptual content is abstract, *then* he has an argument that perceptual content is abstract. But, as far as I can see, Tye’s argument for the identity thesis tacitly appeals to the premise that perceptual content is abstract.

Moreover, Tye gives no *other* argument that perceptual content is abstract. And one is required, because the claim is hardly intuitively correct: if the content of belief can be object-dependent, why can’t the content of perception? In fact, on one of the most sophisticated theories, namely Peacocke’s, perceptual content *is* object-dependent.¹⁷

At the very least, there is no evident reason why the content of perception *couldn’t* be object-dependent (whether or not it actually is). And this suggests an objection. Suppose that Tye is right that the content of our experiences is abstract. Presumably there could be a creature whose experiences were just like ours in content, but with an additional “object-dependent” conjunct. For example, suppose that when one of us looks at a certain tomato (call it ‘ o ’), his visual experience has the content *that (some x) x is red and round*. (We may assume that this content is “nonconceptual”.) Then

¹⁷ According to Peacocke, one “layer” of (“nonconceptual”) perceptual content comprises *protopositions*. Protopositions are simple sorts of Russellian propositions—“A protoposition contains an individual or individuals, together with a property or relation” (1992, 77)—and are therefore not abstract. (For more on Peacocke’s theory, see section 6 below.)

For an extended argument (from a position in many respects opposed to Peacocke’s) for, *inter alia*, the conclusion that perceptual content is object-dependent, see Brewer 1999, ch. 2. For some considerations on Tye’s side, see Davies 1996. It would be a distraction to consider these arguments here. Davies, by the way, claims that Peacocke’s protopositional content is *not* object-involving (310). As I understand Peacocke’s official account, this is not correct; Peacocke does note, however, that an object-independent version of protopositional content is a theoretical option (n. 7, 241).

the content of the creature's visual experience when she looks at the tomato would be *that (some x) x is red and round & o is red and round*. However, because the content of the creature's experience is not abstract, the PANIC theory implies that there is nothing it's like for the creature to look at the tomato. And that seems very odd. How could getting *more* information from vision make the lights go out?

However, various easy repairs can be made to the PANIC theory. For example, if we say that propositions P_1 and P_2 are *abstractly equivalent* iff they are the same modulo the representation of particular individuals, then the PANIC theory could be revised thus: 'Phenomenal character Q = the property of having content *abstractly equivalent* to nonconceptual content P and...'. So, although the A-part of the PANIC theory probably has to go, this objection isn't fatal.¹⁸

6. Nonconceptual content

The most troubling objection to the PANIC theory concerns N. To anticipate: two ways of understanding 'nonconceptual content' yield two interpretations of the PANIC theory (the "state" interpretation and the "content" interpretation), and two corresponding horns of a dilemma. On the state interpretation, arguably experiences do have "nonconceptual content", but the PANIC theory is (at the very least) unmotivated. On the content interpretation, the chief difficulty is that the PANIC theory is seriously underdescribed.

¹⁸ Admittedly, if perceptual content is abstract, then this neatly finesses the problem for the object-dependent view posed by hallucinations, where there is apparently no appropriate object to figure in the content of the experience (cf. 2000, 62). But this isn't a convincing argument unless the problem cannot be solved in other ways. The analogous problem in the philosophy of language is of course the problem of empty names, with Tye's abstractness proposal analogous to the description theory of names. And although the description theory of names *does* neatly finesse the problem of empty names, it is not the only viable solution.

The *Ten Problems* definition of nonconceptual content is quoted in section 3 above; the definition in *Consciousness, Color, and Content* is a little more expansive: “to say that a mental content is nonconceptual is to say that its subject need not possess any of the concepts that we, as theorists, exercise when we state the correctness conditions for that content” (2000, 62).

This needs to be unpacked rather slowly. Start with ‘correctness conditions’. To state the correctness conditions for a content—that is, a proposition—P is simply to specify P using a that-clause: that there is a blue triangle before one, for example. ‘Possessing the concept *F*’ is a little trickier, but I think a close enough approximation to Tye’s usage is this: a subject possesses the concept *F* iff she believes that...*F*...¹⁹ So, for example, if a subject believes that cranberries are red, or that cranberries are not red, or that everything red is colored, then she possesses the concept *red*. And if she possesses the concept *red* then she has *some* belief whose content can be specified using the English word ‘red’.

Next, ‘possessing/exercising the concept *F*’. When we theorists state that the proposition P is the proposition that there is something red and round, we are “exercising” our concepts *red* and *round*. (Note that on this way of explaining “concept” talk, one might regard apparent reference to “the concept *red*”, “the concept *round*”, etc., as a mere *façon de parler*, to be “paraphrased away”; as we will see shortly, this is not Tye’s view.)

¹⁹ See 1995, 108, where Tye mentions that “[h]aving the concept *F* requires, on some accounts, having the ability to use the linguistic term ‘F’ correctly. On other accounts, concept possession requires the ability to represent in thought and belief that something falls under the concept”. He does not officially adopt either of these two kinds of account, but since he thinks non-human animals have concepts (2000, ch. 8), it’s clear that his sympathies lie with the second. And, I think, on the intended construal of ‘the ability to represent...’ the second kind of account is more-or-less equivalent to the one suggested in the text.

Finally, ‘its subject’. Clearly the “subject” of a mental content *P* is supposed to be someone who is in a mental state *S* with the content *P*. So, if Smith believes/hopes/desires that there is something red and round, then Smith is the subject of the content that there is something red and round.

Given this explanation, the conceptual/nonconceptual distinction is most naturally thought of as applying in the first instance to *states*, not to *contents*. And in *Ten Problems* the distinction is first introduced as applying to states: “...perceptual sensations feed into the conceptual system, without themselves being a part of that system. They are nondoxastic or nonconceptual states” (1995, 104²⁰). An abbreviation will be useful: let us say that the concept *F* characterizes the proposition *P* iff *P* = that...*F*... Then (the present version of) the nonconceptual/conceptual distinction can be explained as follows:

Mental state *S* with content *P* is nonconceptual iff someone who is in *S* need not possess any of the concepts that characterize *P*.²¹

We can speak derivatively of nonconceptual *content*: a proposition *P* is nonconceptual iff it is the content of some nonconceptual state. But notice that this account does not imply that “nonconceptual content” is a *special kind* of content. If perceptual experience has nonconceptual content in this sense, the propositions that are the contents of perception might well be perfectly familiar propositions, of the sort that are the contents of belief (Russellian, Fregean, Lewis-Stalnakerian, whatever).

²⁰ See also 1995, 108. Similarly, in *Color, Consciousness and Content* the distinction is first introduced as applying to *experiences*: “experiences of sounds...admit of many more fine-grained distinctions than our stored representations of sounds in memory. Experiences of shapes are likewise nonconceptual” (11).

²¹ Cf. Crane 1992, 143.

Let us call this conception of nonconceptual content the *state* conception. On the state conception beliefs and thoughts are automatically conceptual states; what is controversial is whether perceptual experiences are nonconceptual states—according to the state *view*, they are.

On the state conception, the phrase ‘nonconceptual content’ is somewhat unfortunate, as it suggests a special kind of content. However, according to most theorists of nonconceptual content, the phrase isn’t at all misleading, because it really is a special kind of content. On this alternative conception—the *content* conception—a proposition is nonconceptual iff it isn’t a Fregean Thought—that is, if it isn’t a proposition with Fregean senses or “concepts” (in one sense of the term) as constituents. According to the content *view*, (a) the content of belief and thought is conceptual (i.e. Fregean), and (b) the content of perception is nonconceptual.²² (The useful “state/content view” terminology is taken from Heck 2000.) For example, on Peacocke’s recent proposal, the nonconceptual content of experience is a combination of “scenario content” and “protopositional content”. These abstract objects are built to Russellian specifications: a protopositional content *is* a simple sort of Russellian proposition, while a scenario content is something more complicated, but likewise constructed from materials at the level of reference (Peacocke 1992, ch. 3). The contents of belief and thought, on the other hand, are exclusively conceptual.²³

²² Strictly speaking, (b) should be: the content of perception is *at least partly* nonconceptual. (See, e.g., Peacocke 1992, 88.) This complication will be ignored. See also footnotes 10 and 11 above.

²³ The state and content views are, if not positively muddled up, at least not properly separated in much of the literature (as is pointed out in Stalnaker 1998a, 1998b). A similar commission or omission occasionally infects discussions of narrow content. Sometimes the claim that some content is narrow is simply a claim of *local supervenience*: the property of believing that *p*, for some filling for ‘*p*’, is *intrinsic*. If the belief that *p* has narrow content in this sense, its narrow content is simply the proposition that *p*. And this might well be a perfectly ordinary proposition, of the Russellian, Fregean, or Lewis-Stalnakerian sort, according to taste.

Once this distinction between the state and content views is in place, it is clear that a common argument in the literature—the “richness argument” for nonconceptual content—only supports the *state* view, not the *content* view.²⁴ Tye’s version of the richness argument is this:

Beliefs and thoughts involve the application of concepts. One cannot believe that a given animal is a horse, for example, unless one has the concept *horse*. At a minimum, this demands one has the stored memory representation *horse*, which one brings to bear in an appropriate manner (by, for example, activating the representation and applying it to the sensory input). However...phenomenal seemings or experiences are *not* limited in this way. My experience of red_{19} , for example, is phenomenally different from my experience of red_{21} , even though I have no stored memory representations of these specific hues and hence no such concepts as the concepts red_{21} and red_{19} . These points generalize to the other senses. Phenomenal character, and hence phenomenal content, on my view, is nonconceptual. (1995, 139; cf. 2000, 61-2)

That is, this first sense of ‘narrow content’ doesn’t mark a distinction among *kinds of contents*. But the second sense does: according to it, the narrow content of a belief is special kind of non-propositional abstract object; for example, Fodor once proposed that narrow content is a function from contexts to propositions.

²⁴ The point to follow is an elaboration of Byrne 1996, 264, n.6. Because the richness argument at best supports the state view, Byrne (1996, 263-4) claimed that focus of dispute in the literature was the state view, not the content view. This was an error. Still, some proponents of “nonconceptual content” hold the state view, in particular Crane, who thinks that “perceptions have contents that can be the contents of beliefs” (1992, 155).

That is, to possess the concept F (i.e. to believe that... F ...) one must have, at least, “the stored memory representation F ”. And because it is possible to have a visual experience of red_{21} , without having “the stored memory representation red_{21} ”, one does not have to possess the concept red_{21} in order to have that visual experience. Therefore, a visual experience of red_{21} is “nonconceptual”, or “has nonconceptual content”.

This argument evidently does not even purport to show that experience has nonconceptual content on the content conception. For all this argument says, a subject’s visual experience might have the content that, say, a certain tomato is red_{19} , where the proposition that the tomato is red_{19} is the very same kind of proposition—a Fregean Thought, perhaps—that she can believe.²⁵

Tye’s official argument for nonconceptual content establishes, at best, the *state* view. But Tye in fact holds the *content* view.²⁶ The textual case for this attribution chiefly rests on the manifest inadequacy of the PANIC theory, with the N part interpreted according to the *state* conception. Section 6.1 explains why. Section 6.2 argues that the PANIC theory interpreted according to the *content* conception has problems of its own.

²⁵ The richness argument is in embryo form in Evans 1982, 229, and 125, n. 9; Evans seems to be arguing for the content view, although this is not entirely clear. (A related argument in Dretske 1981, ch. 6; however, plainly Dretske is arguing for something like the state view.) The richness argument is taken to support the content view by Peacocke (1992, 67-8; 1998; for a more guarded view of the argument, see 2001b) and Heck (2001, 489-90); Heck’s version of the richness argument is discussed below in section 6.2. (Neither Peacocke nor Heck can be convicted of conflating the state and content views—in particular, Heck carefully makes this very distinction.) The argument is opposed by McDowell (1994, 56-60; 1998) on the ground that demonstratives like ‘that shade’ can capture the content of color experience (see also Brewer 1999, 170-4; Kelly 2001). However, McDowell appears to concede that the richness argument provides a *prima facie* consideration in favor of the content view.

²⁶ He confirmed this in correspondence. (For a slight complication—not examined further here—see note 30 below.)

6.1 PANIC: the state interpretation

The PANIC theorist—whether she holds the state or content view—is committed to the claim that all beliefs (thoughts, judgments) lack phenomenal character. This is because, she thinks, no belief has nonconceptual content, and on the PANIC theory nonconceptual content is necessary for phenomenal character. And if the PANIC theorist is to offer any *explanation* of why beliefs in general lack phenomenal character, the fact that they are nonconceptual must do the work. Lack of abstractness won't do it, because some beliefs are abstract. Neither will lack of poisedness—but this claim requires a little defense.

Sometimes Tye seems to claim that if a state is poised then by definition it cannot be in “the belief/desire system” (1995, 104, 142). If so, then no belief can be poised. On a more inclusive construal beliefs *can* be poised: a poised belief is one that is available to make a “direct impact” on desires and/or (other) beliefs.

On the *inclusive* construal of poisedness, lack of poisedness cannot explain why beliefs lack phenomenal character, because some beliefs are poised. So, why not adopt the *exclusive* construal of ‘poised’, on which only states outside the “belief/desire system” can be poised? But then the “explanation” that beliefs lack phenomenal character because they are not poised boils down to the unhelpful claim that beliefs lack phenomenal character because they are inside the “belief/desire system”, i.e. because they are either beliefs or desires. This is unsatisfactory (more will be said about this kind of “explanation” in a moment). So there is nothing to be gained by adopting the exclusive construal.

To repeat: any explanation of why beliefs lack phenomenal character must appeal to the fact that they lack nonconceptual content. However, on the *state* interpretation of the PANIC theory, the “explanation” that beliefs lack phenomenal character because they lack nonconceptual content is just as unsatisfactory as the “explanation” in terms of (the exclusive construal of) poisedness. On the state conception, a state *S* with content *P* is a

nonconceptual state just in case it is possible to be in *S* without “possessing the concepts” that characterize the content of *S*; that is, without having *beliefs* (for instance, the belief *P*) in which those concepts figure (see section 6 above). And it immediately follows from this that no belief is a nonconceptual state. Hence, the explanation of why beliefs lack phenomenal character boils down to the unhelpful claim that beliefs lack phenomenal character because it’s not possible to have a belief without having beliefs.

And this is a problem. According to some, conscious beliefs have phenomenal character.²⁷ The PANIC theory’s claim that all beliefs essentially lack phenomenal character is therefore contentious. And even if introspection convinces us that, as a matter of actual fact, beliefs lack phenomenal character, this might just be a contingent truth. It is not a *datum* that beliefs essentially lack phenomenal character. So, if it’s true, it is the sort of thing a theory of consciousness should be able to *explain*. But we have just seen that the PANIC theory, interpreted on the state conception, offers no explanation at all.

Matters are no better when we ask why some states with content *have* phenomenal character. Consider a standard visual experience as of a ripe tomato, and the conscious belief that (some *x*) *x* is red. (We may suppose, with the PANIC theorist, that the belief lacks phenomenal character.) Why does the experience, unlike the belief, *have* phenomenal character? It cannot be because the experience is *abstract*, for the belief is too. Neither can it be because the experience is “poised”, because (we may suppose) the belief is also poised.²⁸ As before, then, the explanatory burden must be borne by nonconceptual content. The fact that the experience has nonconceptual content must be the crucial phenomenology-maker. On the state interpretation, this amounts to the fact that the subject *need not* possess “matching concepts” in order to enjoy the experience.

²⁷ See, for example, Block 1995, 230; Chalmers 1996, 9-10; Peacocke 1999, 205-6.

²⁸ Adopting the exclusive construal of ‘poised’ would not help, for the reason given earlier.

So, for example, the fundamental explanation of why the experience of red_{19} has phenomenal character appeals, not to the fact that subjects who enjoy this experience actually lack the concept red_{19} , but to the modal fact that the experience *could be* enjoyed by a subject who lacked the concept. That is, the experience of red_{19} has phenomenal character because it *could be* enjoyed by a subject who did not believe anything of the form: that... red_{19} ... It is hardly obvious why a subject's enjoying experience e while lacking certain beliefs is relevant to whether e has phenomenal character, and entirely unobvious why the *possibility* of enjoying e while lacking certain beliefs is relevant.

The PANIC theory on the state interpretation does not give a remotely satisfactory explanation of why perceptual experiences have phenomenal character, or why beliefs lack phenomenal character. Since some such explanation is required if we are to have reason to believe the theory, we have no reason to believe it.

6.2 PANIC: the content interpretation

On the content *conception*, nonconceptual content is content that is not *conceptual* or *Fregean*; that is, content that is not composed of “concepts” or Fregean senses. Russellian, Lewis-Stalnakerian, and Peacockean (scenario) contents are consequently examples of (this conception of) nonconceptual content. The PANIC theory interpreted according to the content conception implies the content *view*: beliefs (thoughts, judgments) have conceptual content, and perceptions have nonconceptual content.

A proponent of the content view has a couple of reasons to hold that *linguistic* content—the content of (natural language) sentences, relative to particular contexts of utterance—is also Fregean. First, the traditional route (i.e. Frege's) to the conclusion that the content of belief is Fregean proceeds by establishing first that linguistic content is Fregean. Second, the conclusion that linguistic content is Fregean follows from the premise that belief content is Fregean together with the very plausible premise that the

content of any sentence can be the content of belief (see Peacocke 2001a, 243).²⁹ And, indeed, proponents of the content view invariably endorse the claim that linguistic content is also Fregean.³⁰

Now, although it might be that the PANIC theory supplemented with a well-worked out version of the content view *can* explain why beliefs lack phenomenal character, and why perceptual experiences have it, the immediate problem is that Tye has supplied no good reason in favor of the content view. It is advisable, then, to canvass some other arguments.³¹

Two recent examples are instructive: Heck's version of the richness argument, and Peacocke's discussion of "the most fundamental reasons for acknowledging nonconceptual representational content" (2001b, 613).

²⁹ This premise needs some refinement, because arguably some sentences express propositions that cannot be believed (for example, perhaps no one could really believe that nothing exists).

³⁰ Tye is a Fregean (of the kind who thinks that objects and properties, as well as senses or modes of presentation, are constituents of propositions) (2000, 18). However, he thinks that in some cases beliefs contents can have objects or properties as constituents, with no corresponding modes of presentation: in the special case of "phenomenal concepts", they "refer directly...There is no separate guise that the referent takes in the thinker's thought" (2001, 695; cf. 2000, 28).

Fregeanism, by the way, has been deliberately left at a vague and impressionistic level in this paper, because different theorists understand it differently. For the record, my own sympathies are with a Russellian account (of linguistic content, at any rate).

³¹ According to Tye, the representational vehicles of experiences have a "topographic or maplike structure" (1995, 121; cf. 2000, 70-4), unlike the representational vehicles of beliefs, which have a sentencelike structure (1995, 100) (so Tye thinks there is a language of thought, although not a language of experience). One might try to argue from these differences in *representational vehicles* to a difference in the kinds of *contents represented*. However, Tye does not supply any such argument, and there is no indication that he thinks one could be supplied. Moreover, it would be a confusion to think that a difference in representational vehicles *entailed* a difference in contents represented.

First, Heck's version of the richness argument:

Consider your current perceptual state—and now imagine what a complete description of the way the world appears to you at this moment might be like. Surely a thousand words would hardly begin to do the job. And it is not just that the description would be long: Rather, it seems hard to imagine that your perceptual state, as it is now, has any specific articulation corresponding to the conceptual articulation of a particular one of the many different Thoughts that might capture its content; and it seems at least as hard to imagine that you now possess all the concepts that would be expressed by the words occurring in such a description, even if one could be framed. Before me now, for example, are arranged various objects with various shapes and colors, of which, it might seem, I have no concept. My desk exhibits a whole host of shades of brown, for which I have no names. The speakers to the sides of my computer are not quite flat, but have curved faces; I could not begin to describe their shape in anything like adequate terms. The leaves on the trees outside my window are fluttering back and forth, randomly, as it seems to me, as the wind passes over them.—Yet my experience of these things represents them far more precisely than that, far more distinctively, it would seem, than any other characterization I could hope to formulate, for myself or for others, in terms of the concepts I presently possess. The problem is not lack of time, but lack of descriptive resources, that is, lack of the appropriate concepts (2001, 489-90).

The conclusion of this argument is supposed to be that “the content of perceptual states is different in kind from that of cognitive states like belief” (485). Given the assumption (implicit in the quotation), that the content of belief is conceptual, the content view follows: the content of belief is conceptual, and the content of belief is nonconceptual.

Heck's version of the richness argument overlaps with Tye's: like Tye, he claims that experience represents, say, shades of color "of which, it might seem, I have no concept". For example, one can have an experience of brown₂₇, without having the concept *brown*₂₇. And, as emphasized earlier, this does not have any tendency to show that perceptual experiences have a special kind of content.

However, the quoted passage contains another strand of argument, apparently leading to the conclusion that the content of perception cannot be fully expressed in *any language*—that perceptual content is not *linguistic* content. And if we add the premise that *belief* content can always be fully expressed in language, and further assume that linguistic content is conceptual, then the content view follows. So let us pursue this other strand for a moment.

The claim that perceptual content is not linguistic is not merely the claim that a particular perceiver might lack the vocabulary to express the content of his experiences. This weak claim is no doubt true, but it evidently does not show that perceptual content resists expression in any language, and so does not show that the content of perception and the content of language are different in kind. Hence, Heck's observation that his "desk exhibits a whole host of shades of brown, for which [he has] no names" does not support the claim that perceptual content is not linguistic: presumably the apparent shades of Heck's desk can be captured linguistically with the aid of a paint catalogue. Rather, the crucial consideration is this: "it seems hard to imagine that your perceptual state, as it is now, has any specific articulation corresponding to the conceptual articulation of a particular one of the many different Thoughts that might capture its content" (clearly, given the context, we could replace 'Thoughts' by 'sentences'). The idea here appears to be the *reverse* of the official richness argument. It is not that *perception* is too fine-grained to be captured by the net of language, but rather that *language* is too fine-grained: to attempt to express perceptual content in language inevitably imposes on it a structure that it does not have. So, perhaps, choosing 'p & q' to express the content of an

experience gives it an unwanted conjunctive structure, and other unwanted structures would be introduced by any logically equivalent sentence (say, ‘ $\sim(\sim p \vee \sim q)$ ’).

This is certainly suggestive, but (at any rate in my brief exposition of the point) it is far too slender and elusive a reed to support any weight. Moreover, a similar point about belief would seem to be *equally* suggestive. Extruding beliefs through the templates of language often seems to impose on them unnecessary structure and precision. You realize you have forgotten your car keys, and so go back to the house to pick them up. The fact that you had *some* belief about the keys, together with an appropriate desire, explains your action. But what sentence expresses this belief? There are innumerable candidates: ‘I left the car keys on the kitchen table’; ‘I left the keys on the table in the kitchen’; ‘I forgot to pick up the keys from the table’; ‘The keys are where I left them, on the table’, etc. You are disposed to assent to all of these sentences, and so in this “dispositional” sense you believe the (different) propositions they express, but presumably not all of these beliefs causally explain your behavior. As Dennett puts it, “our linguistic environment is forever forcing us to give—or concede—precise verbal expression to convictions that lack the hard edges verbalization endows them with” (1981, 21). So, although Heck’s second strand of argument hints that *perceptual* content is not linguistic, a parallel strand hints that *belief* content is not (wholly) linguistic either. And this is of course inconsistent with the content view. Nonetheless, I think Heck is onto something here; the issue is examined further in section 7.

Second, Peacocke’s argument for the content view:

Nonconceptual content has been recruited for many purposes. In my view the most fundamental reason—the one on which other reasons must rely if the conceptualist presses hard—lies in the need to describe correctly the overlap between human perception and that of some of the nonlinguistic animals. While being reluctant to attribute concepts to the lower animals, many of us would also

want to insist that the property of (say) representing a flat brown surface as being at a certain distance from one can be common to the perceptions of humans and of lower animals. The overlap of content is not just a matter of analogy, of mere quasi-subjectivity in the animal case. It is literally the same representational property that the two experiences possess, even if the human experience also has richer representational contents in addition. If the lower animals do not have states with conceptual content, but some of their perceptual states have contents in common with human perceptions, it follows that some perceptual representational content is nonconceptual (2001b, 613-4).³²

This argument may be set out as follows:

1. Humans do, and the lower animals do not, “possess concepts”.

Therefore:

2. Humans are in states (e.g. beliefs) with conceptual content, and the lower animals are not.

But:

3. Some of the perceptual states of the lower animals have contents in common with human perceptual states.

Therefore:

³² Cf. Bermúdez 1998, chs. 3 and 4; Evans 1982, 124; and McGinn 1989, 62. McDowell opposes this argument by denying premise (3): “We do not need to say that we have what mere animals have, non-conceptual content, and we have something else as well, since we can conceptualize that content and they cannot. Instead we can say that we have what mere animals have, perceptual sensitivity to features of our environment, but we have it in a special form. Our perceptual sensitivity to our environment is taken up into the ambit of the faculty of spontaneity, which is what distinguishes us from them” (1994, 64). McDowell’s response is also endorsed by Brewer (1999, 177-9).

4. Human perceptual states have a kind of content that is not conceptual, i.e. they have nonconceptual content.

Since, by (2), human belief states have *conceptual* content:

5. The content view is true.

Because ‘possess concepts’ can be glossed in multiple ways, premise (1) can sustain a variety of interpretations. It will be useful to distinguish three of them:

(1*) Humans have beliefs, and the lower animals do not.

(1**) Humans have beliefs with Fregean Thoughts as contents, and the lower animals do not have beliefs.

(1***) Humans have beliefs with Fregean Thoughts as contents, and the lower animals, although they may have beliefs, do not have beliefs with Fregean Thoughts as contents.

How does the argument fare on each of the three corresponding interpretations?

Not well on the first interpretation ((1) = (1*)). (1*) does not support the view that beliefs (unlike perceptions) have a special kind of content, and so does not support (2).

The second interpretation apparently conforms best to Peacocke’s intentions.³³ It holds out more promise of supporting (2), but more needs to be said. On the face of it,

³³ “I shall be taking it that conceptual content is content of a kind that can be the content of judgment and belief. Concepts are constituents of those intentional contents which can be the complete, truth-evaluable contents of judgment and belief” (2001a, 243). And: “concepts...are at the level of Frege’s senses” (1992, 3).

one might reasonably hold (1**) together with the view that perceptual content, in humans and lower animals, is Fregean (i.e. conceptual)—thus denying (2).

On the third interpretation of (1), the lower animals might have beliefs with contents that are not conceptual. And, especially because the focus of the argument is on the *overlap* between humans and the lower animals, perhaps some *human* beliefs have such nonconceptual contents (why not?). So (1), on this interpretation, is in some tension with the conclusion of the argument, because the content view is at least committed to the claim that human belief exclusively has conceptual content. Further, the problem noted for the second interpretation also arises for the third.

Even if the problem for the second interpretation noted above can be overcome, there is the additional difficulty of justifying the claim that “the lower animals” (which Peacocke takes to include cats and dogs, and perhaps monkeys and apes³⁴) enjoy perceptual experiences with contents in common with human perceptual experiences, while lacking beliefs. These issues are too large to be discussed here, but once it is conceded that having beliefs is not constitutively tied to speaking a language (as Peacocke himself is at pains to emphasize), then surely the burden of proof is on those who *deny* that humans and the lower animals have beliefs in the same robust sense.

Peacocke’s line of argument for the content view is an uphill struggle. What’s more, Tye himself would reject it completely. For according to him, *fish* have beliefs, and possess concepts (2000, 176-7). Notice that if fish *lack* beliefs, then none of their states

³⁴ “Cats, dogs, and animals of many other species, as well as human infants, perceive the world, even though their conceptual repertoire is limited, and perhaps even nonexistent...the “soft line”...says that some of the conscious perceptual states with representational content enjoyed by mature humans can be enjoyed by nonlinguistic animals without concepts, or with only minimal conceptual capacities” (2001a, 260). And: “the soft line is right” (261). (I have ignored the hedging about “minimal conceptual capacities”. It is absent in Peacocke’s 2001b, and so presumably Peacocke does not regard it as particularly significant.)

are poised: no state “stands ready and available” to affect beliefs.³⁵ So, if fish lack beliefs, then the PANIC theory implies that there is nothing it’s like to be a guppy. Guppy consciousness is no doubt a bit fishy, but it is almost universally (and rightly) held that dogs and apes are phenomenally conscious. Hence, any reasonable PANIC theorist is committed to the view that these animals have beliefs, which puts him on a collision course with Peacocke’s “fundamental reason” for nonconceptual content.

To sum up the discussion of nonconceptual content. The PANIC theory interpreted on the *state* conception of nonconceptual content is inadequate (as Tye would no doubt agree). The right interpretation builds the *content view* into the PANIC theory. However, we have found no reason to believe the content view: that beliefs have conceptual content and perceptual experiences have another kind of content—nonconceptual content. Further, even if perceptual content is nonconceptual, Tye does not give any positive account of it. Lastly, because of the previous point, it is completely obscure why nonconceptual content (on the content conception) is part of ingredient X.

7 $X=P+A+N$ revisited

If the argument so far is correct, Tye has misidentified ingredient X: it is not P+A+N. However, there are some important insights underlying his proposal—specifically the selection of P and N.

First, P. Its main role in the PANIC theory is to account for blindsight. In blindsight, the subject has a quasi-experience, say as of an ‘O’ before her, but (it is natural to say) *she herself* is unaware, or not conscious, that there is an ‘O’ before her.

³⁵ “Standing ready and available” to affect *desires* is also sufficient for poisedness; but we may fairly assume that desire and belief go together: lacking one entails lacking the other.

What I take to be the basic intentionalist insight about blindsight is this. The missing ingredient is not a non-intentional quale, or even a special kind of content, but simply *the conscious subject herself*. It does not seem *to her* that there is an ‘O’ before her. Assuming for simplicity that the content of her quasi-experience is the proposition *that there is an ‘O’ before her*, all that is required for phenomenal character is that it seems *to the subject* that there is an ‘O’ before her.

Although this may be an insight, it is not of much help in furthering reductive or physicalistic ambitions. However, if one adopts some sort of Humean bundle-theory of the self, as I suspect Tye tacitly does, then the problematic notion of the conscious subject herself may be cashed out in terms of certain privileged mental states. Specifically, in Tye’s theory, it’s seeming to the subject that *p* is reduced to the self-free fact that a state with the content that *p* “stands ready and available to make a direct impact on beliefs/desires”. As we have seen, this does not seem to work. But the fundamental problem is with Tye’s reductive ambitions, not with the basic insight about blindsight.

Second, N. Here Tye’s insight is that a theory of consciousness *does* need a special kind of content. Nonconceptual content, though, is the wrong candidate. It is supposed to be content that cannot be *believed* (and therefore cannot be linguistically expressed). What we want instead is content that *can* be believed, but that *cannot* be linguistically expressed. (See again the discussion of Heck in section 6.2.)

I shall now outline an argument for this claim, based on Jackson’s (1982) knowledge argument together with a perceptive remark of Lewis’s.³⁶ Assume, first, that

³⁶ For another way of approaching the same conclusion, see Byrne forthcoming; an important related discussion is in Thau forthcoming.

knowing what it's like to enjoy an experience is propositional knowledge.³⁷ When black-and-white Mary sees a ripe tomato for the first time, and thereby comes to know what it's like to see red, she comes to know some proposition. If one were forced to choose a sentence to express this proposition, a plausible candidate would be 'Seeing red is like *this*', where we imagine Mary uttering this sentence while looking at a tomato. So, assuming for the moment that the proposition Mary learns is linguistically expressible, we may write it thus:

(M₁) Seeing red is like *this*.

Essentially the same piece of knowledge can be put in helpful jargon as follows:

(M₂) Having an experience that represents objects as red is like *this*.

For an intentionalist like Tye, Mary comes to know M₂, not by directly introspecting her *experience*, but by attending to the colors in the scene before her eyes: "Our attention goes outside in the visual case, for example, not to the experience *inside* our heads. We attend to one thing—the external surface and qualities—and yet *thereby* we are aware of something else, the 'feel' of our experience" (2000, 51-2).³⁸ In other words, Mary is in a position to know M₂ once she knows:

(M₃) An experience that represents objects as red represents them like *this*.

³⁷ See, for example, Lycan 1996, ch. 5; a closely related claim, that "knowing how" is a species of "knowing that", is argued for in Stanley and Williamson 2001. Tye himself holds that "knowing what it is like is best captured by a disjunction of introspective knowing-that and knowing-how" (2000, 16).

³⁸ See also Byrne 2001; Dretske 1995, ch. 2; Shoemaker 1994.

Note that M_3 is a proposition that specifies the distinctive way red objects are represented in visual experience; that is, it specifies the *content* distinctive of experiences as of red objects. (Of course, an anti-intentionalist would deny that knowing M_3 puts Mary in a position to know what it's like to see red.)

Now to Lewis's perceptive remark: "Our intuitive starting point wasn't just that *physics* lessons couldn't help the inexperienced to know what it's like. It was that *lessons* couldn't help" (1988, 281). Therefore, since knowing M_3 would help imprisoned Mary to know what it's like, the proposition M_3 cannot be taught by a *lesson*.

But what is a "lesson"? In one sense, showing Mary a ripe tomato is giving her a lesson, but obviously that is not what Lewis means. Instead, it's clear that he means *linguistic* lessons. No matter how many books imprisoned Mary reads, and lectures she hears, she won't come to know what it's like to see red. And this is not because there are some sentences that Mary can't understand. Although she hasn't had the experience of seeing red objects, that does not prevent her from understanding any linguistic expression (so, for example, she can understand the word 'red' while imprisoned). Of course, there will be uses of demonstratives that could not occur in lessons Mary has while imprisoned, in particular an utterance of 'An experience as of red objects represents them like *this*' in the presence of a tomato. And such an utterance of that sentence expresses—we have been supposing—the proposition M_3 . But this does not mean that the proposition M_3 —if it really is expressed by that sentence—could not be taught to imprisoned Mary. Plausibly, any proposition expressed using a demonstrative could be expressed in a demonstrative-free way: for example, the proposition expressed by 'That man is drinking a martini' (pointing at Tye) is arguably expressed by the demonstrative-free sentence 'Tye is drinking a martini'. Assume this is correct. Then, if M_3 really is expressed by an appropriate utterance of 'An experience that represents objects as red represents them like

this', we could teach M_3 to imprisoned Mary: no demonstration of ripe tomatoes is needed.

All the premises are now in place (albeit with minimal defense). If M_3 can be linguistically expressed, then Mary can know M_3 while imprisoned, and thereby know what seeing red is like. But she can't know this while imprisoned. Therefore M_3 can't be linguistically expressed. Our supposition that M_3 is expressible using a demonstrative is a ladder that must be kicked away: in using a demonstrative, we were trying to say what can't be said. We can, however, *communicate* or *convey* M_3 , by uttering the sentence 'An experience that represents objects as red represents them like *this*' in the presence of a ripe tomato; at least, M_3 can be communicated in this way to those who have the appropriate sort of experience. (And, I presume, I have succeeded in communicating M_3 to *you*.) For familiar Gricean reasons, a proposition can be communicated by uttering a sentence in a context, even if the proposition is not the semantic content of that sentence relative to that context. Hence, it doesn't follow from the fact that M_3 can be communicated by uttering a sentence in a context, that M_3 is the semantic content of that sentence relative to that context; neither does it follow that M_3 is the semantic content of *some* sentence.³⁹

In other words: knowing linguistically expressible propositions is *not* sufficient for knowing what it's like, but knowing propositions that specify the content of perception is. Hence, the content of perception cannot be completely expressed in language. The limits of my language aren't the limits of my world, after all.

³⁹ What *is* the proposition expressed by 'An experience as of red objects represents them like *this*' (uttered in the appropriate context)? Arguably, it is the trivial proposition that an experience that represents objects as red represents them as red. That is certainly something that Mary could know while imprisoned.

Assuming that the gaps in this argument can be filled, we need a positive account of both linguistic and perceptual content. And here Peacocke's work on nonconceptual content at least provides a model of how to proceed.

That completes our investigation of the PANIC theory; I hope the theory's virtues, and the difficulty of the problems it sets out to solve, were exhibited along the way. The provisional conclusion is that ingredient X is a certain kind of non-linguistic content plus the subject of experience. This does not deserve to be called a theory of phenomenal consciousness—but perhaps it is a signpost pointing in the right direction.

References

- Bermúdez, J. L. 1998. *The Paradox of Self-Consciousness*. Cambridge, MA: MIT Press.
- Block, N. 1990. Inverted earth. *Philosophical Perspectives* 4, ed. J. Tomberlin. Atascadero, CA: Ridgeview.
- Block, N. 1995. On a confusion about a function of consciousness. *Behavioral and Brain Sciences* 18: 225-47.
- Block, N. Forthcoming. Mental paint. *Others on Burge: 10 Essays With Responses From Tyler Burge*, ed. M. Hahn and B. Ramberg. Cambridge, MA: MIT Press.
(Available at:
www.nyu.edu/gsas/dept/philo/faculty/block/papers/mentalpaint.pdf.)
- Braddon-Mitchell, D., and F. Jackson. 1996. *Philosophy of Mind and Cognition*. Oxford: Blackwell.
- Brewer, B. 1999. *Perception and Reason*. Oxford: Oxford University Press.
- Burge, T. Forthcoming. Reply to Block. *Others on Burge: 10 Essays With Responses From Tyler Burge*, ed. M. Hahn and B. Ramberg. Cambridge, MA: MIT Press.
- Byrne, A. 1996. Spin control: comment on John McDowell's *Mind and World*. *Philosophical Issues* 7, ed. E. Villanueva. Atascadero, CA: Ridgeview.
- Byrne, A. 2001. Intentionalism defended. *Philosophical Review* 110: 199-240.
- Byrne, A. Forthcoming. Something about Mary. *Grazer Philosophische Studien* (special issue on the philosophy of Terry Horgan).
- Carruthers, P. 2000. *Phenomenal Consciousness*. Cambridge: Cambridge University Press.
- Chalmers, D. J. 1996. *The Conscious Mind*. Oxford: Oxford University Press.

- Crane, T. 1992. The nonconceptual content of experience. *The Contents of Experience*, ed. T. Crane. Cambridge: Cambridge University Press.
- Davies, M. 1996. Externalism and experience. *Philosophy and Cognitive Science: Categories, Consciousness, and Reasoning*, ed. A. Clark, J. Ezquerro, and J. M. Larrazabal. Dordrecht: Kluwer. Reference to the reprinting in *The Nature of Consciousness*, ed. N. Block, O. Flanagan, and G. Güzeldere, Cambridge, MA: MIT Press, 1997.
- DeLucia, P. R., and J. Hochberg. 1991. Geometrical illusions in solid objects under ordinary viewing conditions. *Perception and Psychophysics* 50: 547-54.
- Dennett, D. C. 1981. True believers. Reprinted in Dennett, *The Intentional Stance*, Cambridge, MA: MIT Press, 1987.
- Dretske, F. 1981. *Knowledge and the Flow of Information*. Oxford: Blackwell.
- Dretske, F. 1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Evans, G. 1982. *The Varieties of Reference*. Oxford: Oxford University Press.
- Fodor, J. 1990. *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.
- Heck, R. G. 2000. Nonconceptual content and the “space of reasons”. *Philosophical Review* 109: 483-523.
- Kelly, S. D. 2001. The non-conceptual content of perceptual experience: situation dependence and fineness of grain. *Philosophy and Phenomenological Research* 62: 601-8.
- Kirk, R. 1994. *Raw Feeling*. Oxford: Oxford University Press.
- Levine, J. 2001. *Purple Haze*. Oxford: Oxford University Press.
- Lewis, D. K. 1988. What experience teaches. Reprinted in Lewis, *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press, 1999.

- Lycan, W. G. 1996. *Consciousness and Experience*. Cambridge, MA: MIT Press.
- McDowell, J. 1994. *Mind and World*. Cambridge, MA: Harvard University Press.
- McDowell, J. 1998. Reply to commentators. *Philosophy and Phenomenological Research* 58: 403-31.
- McGinn, C. 1989. *Mental Content*. Oxford: Basil Blackwell.
- McGinn, C. 1991. *The Problem of Consciousness*. Oxford: Blackwell.
- Metstre, D. R., et al. 1992. Perception of optical flow in cortical blindness: a case report. *Neuropsychologia* 30: 783-95.
- Milner, A. D., and M. A. Goodale. 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.
- Peacocke, C. 1983. *Sense and Content*. Oxford: Oxford University Press.
- Peacocke, C. 1992. *A Study of Concepts*. Cambridge, MA: MIT Press.
- Peacocke, C. 1998. Nonconceptual content defended. *Philosophy and Phenomenological Research* 58: 381-8.
- Peacocke, C. 1999. *Being Known*. Oxford: Oxford University Press.
- Peacocke, C. 2001a. Does perception have a nonconceptual content? *Journal of Philosophy* 98: 239-64.
- Peacocke, C. 2001b. Phenomenology and nonconceptual content. *Philosophy and Phenomenological Research* 62: 609-15.
- Shoemaker, S. 1994. Self-knowledge and 'inner sense'. Lecture III: the phenomenal character of experience. *Philosophy and Phenomenological Research* 54: 219-314.
- Stalnaker, R. 1984. *Inquiry*. Cambridge, MA: MIT Press.

- Stalnaker, R. 1998a. What might nonconceptual content be? *Philosophical Issues* 9, ed. E. Villanueva. Atascadero, CA: Ridgeview.
- Stalnaker, R. 1998b. Replies to comments. *Philosophical Issues* 9, ed. E. Villanueva. Atascadero, CA: Ridgeview.
- Stampe, D. 1977. Towards a causal theory of linguistic representation. *Midwest Studies in Philosophy* 2: 42-63.
- Stanley, J., and T. Williamson. 2001. Knowing how. *Journal of Philosophy* 98: 411-44.
- Steward, H. 1997. *The Ontology of Mind*. Oxford: Oxford University Press.
- Thau, M. Forthcoming. *Consciousness and Cognition*. Oxford: Oxford University Press.
- Tye, M. 1995. *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
- Tye, M. 2000. *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.
- Tye, M. 2001. Oh yes it is. *Mind* 110: 695-7.