Précis of:
Consciousness, Color, & Content (MIT Press, 2000)

by Michael Tye

Department of Philosophy
The University of Texas at Austin
Austin, Texas 78750
USA

In 1995, in my book, *Ten Problems of Consciousness* (Bradley Books, MIT Press), I proposed a version of the theory of phenomenal consciousness now known as *representationalism*. The present book, in part, consists of a further development of that theory along with replies to common objections. It is also concerned with two prominent challenges for *any* reductive theory of consciousness: the explanatory gap and the knowledge argument. In addition, it connects representationalism with two more general issues: the nature of color and the location of the phylogenetic dividing line between those creatures that are phenomenally conscious and those that are not.

The book, which is made up of eight essays integrated into a whole, is divided into three parts. Part I focuses upon the explanatory gap and the knowledge argument. It aims to show that the right general strategy for dealing with these objections to reductionist theories of consciousness is to hold that the concepts deployed when subjects introspect their experiences and form a conception of their phenomenal character — phenomenal concepts, as I call them — are conceptually irreducible. A theory is developed of phenomenal concepts, one consequence of which is that questions posing the supposed explanatory gap are questions that cannot *possibly* be answered. They are thus not genuine questions and the claim that there are answers to these questions that the reductionist fails to provide is seen to be a kind of cognitive illusion.

Part II, which consists of four essays, is devoted to representationalism itself. It opens with a summary of representationalism and its motivations. Particular attention is paid to the development of the so-called "transparency intuition" on behalf of the theory. The following three chapters deal with objections to representationalism that take the form of putative counter-examples.

The first class of these consists of actual, real-world cases in which, it is claimed, perceptual experiences are the same representationally but different phenomenally. These are the focus of Chapter 4. Another class consists of imaginary cases in which supposedly experiences are identical representationally but inverted phenomenally.

Symposium on Michael Tye's *Consciousness, Color and Content* : 1-2.
<http://host.uniroma3.it/progetti/kant/field/tyesymp.pdf>

These cases, along with a modified representational theory proposed by Sydney Shoemaker, are the focus of Chapter 5. A third class of putative counter-examples is made up of problem cases in which allegedly experiences have different representational contents (of the relevant sort) but the same phenomenal character. Ned Block's Inverted Earth example (*Philosophical Perspectives* 4, 1990) is of this type. Counter-examples are also sometimes given in which supposedly experience of one sort or another is present but in which there is no state with representational content. Swampman -- the molecule by molecule replica of a notable philosopher (Donald Davidson), formed accidentally by the chemical reaction that occurs in a swamp when a partially submerged log is hit by lightning -- is one such counter-example, according to some philosophers. Chapter 6 presents replies both to the Inverted Earth example and to Swampman.

Part III of the book deals with some more general issues, one of which is potentially threatening to representationalism and the other of which representationalism enables us to make progress upon. The potential threat is posed by color (and other so-called "secondary qualities"). For reasons which are made clear in Chapters 3-6, representationalism of the sort I endorse requires an objectivist account of color. It does not require that colors be *external*, objective entities, but this is certainly the view of color that goes most naturally with representationalism. This is also, I believe, the commonsense view of color. Unfortunately, according to many color scientists and some philosophers, colors cannot be objective entities of the sort commonsense supposes. Commonsense supposedly conflicts with modern science on color, and commonsense supposedly has no way of accommodating the distinction between unitary and binary colors. I argue that this is quite wrong. Chapter 7 may thus be seen as a vindication of commonsense and thereby indirectly a defense of representationalism with respect to color.

The view articulated in Chapter 7 of color is developed further in a more recent essay, co-written with Peter Bradley, entitled "Of Colors, Kestrels, Caterpillars, and Leaves," (*Journal of Philosophy*, October 2001). One objection that has been raised to the view presented of color in the book is that it delivers the wrong results for the surface colors of some everyday objects. This objection fails to note that the proposal I make in Chapter 7 is made on the assumption that the oversimplified opponent processing model from color science I appeal to there is correct. As I explicitly say in the chapter, any counter-example to the objectivist proposal I offer for color will also be a counter-example to the oversimplified opponent processing model. Complicate the latter appropriately to handle the counter-example and the former, with corresponding complications, will handle it too. This general claim is illustrated with examples in the *Journal of Philosophy* essay.

The final chapter considers an important question about consciousness on which philosophers have been largely silent, namely: Where, on the phylogenetic scale, does phenomenal consciousness cease? I address this question from the perspective of representationalism and I argue that consciousness extends beyond the realm of vertebrates to such simple creatures as honey bees.

DON'T PANIC: Tye's Intentionalist Theory of Consciousness*

by Alex Byrne

Department of Linguistics and Philosophy
E39-245, MIT
Cambridge, MA 01239
USA

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.

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.

* 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.

Symposium on Michael Tye's *Consciousness, Color and Content* : 3-27.
<http://host.uniroma3.it/progetti/kant/field/tyesymp.pdf>

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 (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

---

[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.

[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.

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":

> 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, *inter*modal intentionalism implies *intra*modal 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, forthcoming)—the answer is yes, while intentionalists disagree.[6] Again, some 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.

---

[5] Lycan (1996) is an example of an *intra*modal 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).

[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

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 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]

---

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.

[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 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*

---

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

[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

the content of my 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

---

> 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.

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.

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

[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.

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

---

[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.

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 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 Carrruthers 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 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.

---

[14] Some caution is needed. The patient was not completely blind, having a small amount of macular and perifovial 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.

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-*in*dependent 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 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 the content of the creature's visual experience when she looks at the tomato would be *that*

---

[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.

[17] According to Peacocke, one "layer" of ("nonconceptual") perceptual content comprises *protopropositions*. Protopropostions are simple sorts of Russellian propositions—"A protoproposition 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 protopropositional 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 protopropositional content is a theoretical option (n. 7, 241).

(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.

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,

---

[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.

[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.

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.)

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, 104[20]). 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).

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

---

[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.

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 "protopropositional content". These abstract objects are built to Russellian specifications: a protopropositional 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]

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, 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

---

[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. 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).

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.

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

---

[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.)

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. 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

---

[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.

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-Stalnakarian, 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]

---

[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

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).

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 brown_{27}, without having the concept $\textit{brown}_{27}$. 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

---

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.

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, '~(~p v ~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:
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,

---

[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).

[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 apes[34]) 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 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.

---

[34] "Cats, dogs, and animals of many other species, as well as human infants, perceive the world, even though there 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.)

[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.

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.

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

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 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:

---

[36] For another way of approaching the same conclusion, see Byrne 2002; an important related discussion is in Thau 2002.

[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).

(M_1) Seeing red is like *this*.

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

(M_2) Having an experience that represents objects as red is like *this*.

For an intentionalist like Tye, Mary comes to know M_2, 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 M_2 once she knows:

(M_3) An experience that represents objects as red represents them like *this*.

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.

---

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

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.

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.

[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.

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.

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: 49-90.

Byrne, A. 2002. Something about Mary. *Grazer Philosophische Studien* 63: 123-40 (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.

Mestre, 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. 2002. *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.

To PANIC or Not to PANIC? - Reply to Byrne

by Michael Tye

Department of Philosophy
The University of Texas at Austin
Austin, Texas 78750
USA

1) Alex Byrne opens his insightful comments by suggesting that the Müller-Lyer illusion provides a straightforward counter-example to my causal co-variational account of the representational content of sensory states. Byrne notes that if optimal 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" (as I say they do, for the case of evolved creatures, in *Ten Problems*), then the Müller-Lyer illusion occurs under optimal conditions. Yet in this case, one's visual experience does not represent that the two lines are of the same length, as the account of content as causal covariation under optimal conditions would seem to predict. Rather, it represents incorrectly that the lines are of different lengths.

In *Consciousness, Color, & Content*, I say a little more about optimal conditions than I did in the earlier book. On the latest development of the causal covariation proposal, optimal conditions are essentially 'no interference' conditions. Given this understanding, my remark quoted above about optimal conditions is best taken to be like the remark that rocks fall to earth when dropped. Typically this is the case, but there are rare exceptions, as when helium balloons are attached to the rocks.

Understood in this way, optimal conditions do not obtain when one views the Müller-Lyer diagram. The standard explanation of the illusion is that, given our experiences with edges and corners in the external environment, the differing arrow heads at the ends of the two parallel vertical lines lead our visual systems to represent, at a subpersonal level, the left hand vertical line as further away than the right (the left hand line is 'taken' in three-dimensional projection to be the edge between two surfaces receding from the subject, the right one is 'taken' to be the edge between two surfaces sloping towards the subject). But the vertical lines are actually of the same length and so their projections on the retina are of the same length too. The difference in distance away of the lines is compatible with the same retinal projection only if one line is longer than the other. Accordingly, our visual systems infer that tne left hand line is longer, and this is what we experience. Obviously this account has it that there is a kind of interference created by the differing directions of the arrow heads in our processing of the distance away of the lines. So, conditions are *not* optimal.

Symposium on Michael Tye's *Consciousness, Color and Content* : 28-33.
<http://host.uniroma3.it/progetti/kant/field/tyesymp.pdf>

2) My PANIC theory of phenomenal character requires of a state with phenomenal character that its content be suitably poised. Byrne's second objection concerns this condition. Before I turn to it, a quick comment on phenomenal character. Byrne says that we should agree at the outset that the phenomenal character of a state is a *property* of that state. I don't see why we should so agree. To be sure, we talk of states *having* phenomenal character, but nothing in ordinary usage or thought commits us to the view that phenomenal characters are properties. After all, we talk of beliefs as having content and of words as having meaning, but we don't take belief contents or word meanings to be properties of beliefs and words respectively. Why do that from the outset for phenomenal character? Buy into the Cartesian view of experiences as inner ideas or pictures viewed by an inner eye and it may be natural to take the 'feel' of an experience as a property of the idea or picture. But that isn't commonsense. It's philosophical dogma — precisely the dogma which representationalism opposes.

So, what is the objection about poisedness? Essentially, Byrne's claim is that the condition isn't strong enough for my purposes. States can have PANIC without having any phenomenal character. In making this charge, Byrne relies on my characterization of poisedness as being the property of standing ready and available to make a direct difference with respect to beliefs and/or desires.

3) I plead guilty to failing to develop the poisedness condition fully enough but not guilty to imposing a condition that is inherently too weak. Let me focus first upon the example of visual experience. Intuitively, visual experiences are not themselves beliefs but they are *apt* for the production of beliefs. Admittedly, some states that might reasonably be classified as visual experiences, for example, seeing that the table is covered with books, already involve beliefs or belief-like states. But such states, in my view, are hybrid, having a visual experience proper and a belief or thought as components.

Visual experiences proper are not apt for the production of any old beliefs, however. Intuitively, each visual experience is the direct basis for the formation of a belief about the perceptible qualities represented by the experience. Each experience is also, in creatures equipped with the capacity to introspect, the direct basis for the formation of beliefs about the experience and its content.

The content of the visual experience proper supplies the input to the relevant belief-forming processes, where the role of the belief-forming processes is to generate beliefs of the sorts just described. But the appropriate beliefs are not always formed, of course; for introspection can malfunction and, at least in the case of external belief formation, other background beliefs can interfere. There is also the possibility that attention is not appropriately directed.

A visual experience has a poised content, then, in my view, so long as it is apt for the production in the *right* ways of the *right* beliefs. In the case of bodily experiences, desire is also relevant. The experience of pain, for example, is the direct basis for the the desire to protect oneself, to avoid damage.

4) We can now see why Byrne's proposed counter-examples fail. A green, circular after-image experience has a poised content since it is apt to produce the belief that one is having such an experience and further it is apt to produce the belief that something green and circular is present (albeit something filmy and floating in space), since it *would* generate the latter belief, *were* there no interference from background beliefs. Going by the experience *alone*, one is inclined to believe that something green and circular is present. In the waterfall illusion, one has an experience of something moving and not moving. This experience has a poised content too. For again it is the direct basis for the appropriate introspective belief. And even though the experience (in virtue of its content) does not cause the belief that something is moving and not moving, as one directs one's attention to the rocks by the side of the waterfall, it would do so, if one were to take the experience at face value without being influenced by other backround beliefs about the impossible.

Mestre's cortically blind patient undergoes visual states whose contents are not poised, however. For he cannot form beliefs about the contents of those states via introspection even though he has the usual human power to introspect and there is nothing wrong with his introspective mechanisms. His visual states are not sufficiently 'close', as it were, to some of the relevant belief-forming processes.

The blindsight subject who undergoes an experience of something red, round, and bulgy, as he views a ripe tomato, and who forms no beliefs about what he sees, is in a visual state that triggers the desire to eat (via, we are to suppose, "some quirk of his inner wiring"). But, contra Byrne, that state's content is not poised, as explained above.

5) Byrne's next objection concerns my abstractness requirement on phenomenal content. He says that "there is no evident reason why the content of perception *couldn't* be object-dependent" and he takes this to show that for possible creatures whose perceptual states are of an object-dependent sort, on the PANIC theory, there is nothing it is like for them to perceive things. This, he claims, is an unacceptable result.

Consider my hallucinating a pink rat. It should be uncontroversial that my perceptual state is delusive. But it cannot be delusive, unless it has an inaccurate representational content. Such a content clearly is not object-dependent, for there is no (relevant) object — I am hallucinating. This perceptual content, then, is abstract.

Since arguments from transparency and the intensionality of phenomenal discourse seem to me to provide strong reasons for thinking that phenomenal character is representational content of some sort and since perceptual states have phenomenal character whether or not their subjects are hallucinating, we have, I suggest, strong reasons for thinking that phenomenal character is representational content of an *abstract* sort, whatever further conditions are necessary.

Of course, this line of argument presupposes that there is no difficulty in holding that abstract perceptual content attaches to states involved in cases of veridical perception. But why shouldn't we hold that? Even if veridical perceptual states have object-dependent contents, it does not follow from this that they don't have abstract contents

too. If I see a picture as a duck, then my visual state has a conceptual content, but it doesn't follow that it lacks any nonconceptual content. There are, it seems to me, many layers of perceptual content[1]; and the possession by a perceptual state of one of these layers does not preclude it from having others.

6) Byrne notes that I hold what he calls "the content view" of nonconceptual content and not "the state view". On the content view, nonconceptual content is content of a sort that cannot be had by beliefs. He goes on to argue that, on this view, "the PANIC theory is seriously underdescribed." But what exactly is the problem? Here is what Byrne says:

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.

First, some remarks on content. In my view, phenomenal 'looks' contexts are intensional to the following extent: in the context 'looks F', where 'looks' is used phenomenally, 'G' can safely be substituted for 'F' so long as 'F' and 'G' pick out the same property.[2] In belief contexts, the intensionality runs deeper. Beliefs use concepts and concepts individuate in a fine-grained way. For example, the concepts WATER and H_2O are different concepts as are the concepts BENJAMIN FRANKLIN and THE INVENTOR OF BIFOCALS, even though the members of each pair co-refer. This is why one can believe that water is thirst quenching without believing that H_2O is thirst, for example. Admittedly, phenomenal concepts, in my view, are a special case. They refer directly.[3] But they are, I believe, the only concepts that so refer. In saying this, I am not implying that indexical concepts, for example, are tantamount to descriptions. The concept I, for example, contributes both its referent and a first-person mode of presentation to the content of any belief that exercises it. But the first-person mode is not descriptive and it does not fix the referent. For this and many other concepts, a two-factor theory of content seems best.

Since one cannot have a belief that uses only phenomenal concepts — at a minimum other logical concepts must be involved — belief content is always at least partly

---

[1] Thus Chris Peacocke (1992, "Scenarios, Concepts, and Perception,' in *The Contents of Experience*, ed. By T. Crane, 105-35, Cambridge, Cambridge University Press) allows scenario contents, protopropositional contents, and conceptual contents. A perceptual state can have an accurate scenario content, for example, but an inaccurate protopropositional one.

[2] I am categorically not making this claim for the context 'looks to be *F*' or the context 'looks as if it is *F*' or indeed for any nonphenomenal 'looks' contexts.

[3] Even so, their identity as *phenomenal* concepts demands that they play the right functional role. This role ensures that phenomenal thought (belief) types play a different role in rationalizing explanations than nonphenomenal thought (belief) types. See here "A Theory of Phenomenal Concepts," *Philosophy* forthcoming.

sensitive to the concepts exercised. Change any nonphenomenal concept or replace a phenomenal concept with a nonphenomenal one and the content of the belief changes. This is not true with content of the sort had by purely phenomenal states. If the picture *looks* square to Samantha, for example, then it *looks* the shape of the tiles, and vice versa, assuming square is the shape of the tiles (though, of course, the picture need not look the shape the tiles look). Moreover, the point extends to nonconceptual content generally. So, on my view, nonconceptual content is not content of the sort that belongs to beliefs.

Byrne says that I have supplied no good reason in favor of the content view. But I just did. And that reason is in *Consciousness, Color, & Content*. Of course, the reason is not one that is likely to persuade Byrne, since he does not accept the orthodox view that belief contexts are intensional in the way I have noted. But the alternative he proposes, under which beliefs have very coarse-grained contents, has well known counter-intuitive consequences.

Beliefs, then, on my view, lack phenomenal character because they lack nonconceptual content. But this is not the only reason for their having no 'feel'. They also are not poised. Beliefs are not conscious states at all. They are manifested in consciousness in thoughts, just as they are manifested in speech by the production of sentences. Beliefs no more have poised contents than do any of the states in early vision (e.g., states representing changes in light intensity). For a state's content to be poised, it must be available to the processes that form beliefs and desires (and further not just any old belief/desire formation processes). And that requires that the state be at the interface of the conceptual and nonconceptual domains, on the nonconceptual side. Beliefs themselves, then, do not play the right role, not even phenomenal beliefs.

Perhaps it will be argued that in tailoring the account of poisedness so as to rule out phenomenal beliefs as the bearers of phenomenal character I have made the proposal ad hoc. I disagree. A priori reflection upon the role that experiences qua experiences play with respect to cognition delivers the poisedness condition. Moreover, there are independent reasons for denying that beliefs have phenomenal character. Indeed there are such reasons for supposing that thoughts, the conscious manifestations of beliefs, do not have phenomenal character.[4] Whatever phenomenal character goes with a thought attaches to associated images, most notably linguistic auditory images. And these images are experiences in their own right.

[4] For some relevant comments here, see my review of Galen Strawson's *Mental Reality*, *Journal of Philosophy*, 1996, 421-424. For the full argument, see my *The Unity of Consciousness*, Chapter 3, in preparation.

Tye on Phenomenal Character and Color

by Barry Maund

Department of Philosophy
The University of Western Australia
35 Stirling Highway
Crawley W. A. 6009
Australia

Tye's principal aim is to defend what he calls 'strong representationalism', a view that aims to tell us precisely what the phenomenal character of our experiences is: it is the same as representational content (that meets further specifiable conditions). More precisely, it is given as the theory that runs by the title PANIC: phenomenal character is one and the same as Poised, Abstract, Nonconceptual, Intentional Content. It is nonconceptual in the sense that the subject need not possess any of the concepts that we, as theorists exercise when we state the correctness conditions for that content. It is abstract in that it is content into which no particular concrete object or surface enters. What is crucial to phenomenal character is explained in terms of the distinctive functional role the experiences and feelings, qua bearers of phenomenal character, play. They arise at the interface of the nonconceptual and conceptual domains and they stand ready and available to make a direct impact on beliefs and desires.

One of the important applications of Tye's Representationalist thesis is to the vexed question of colors and color experiences. The thesis is used to argue that we can explain the phenomenal character of color experiences without admitting any phenomenal, subjective qualities of the experiences (without admitting qualia in the strong sense).

His setting up the framework of the book in the first two chapters is impressive. Tye provides an excellent analysis of Jackson's Knowledge Argument concerning Mary, the all-knowing black and white vision expert. Tye is persuasive in bringing out the strengths and weaknesses of Lewis's Ability Hypothesis (in Ch.1), as an explanation of what Mary discovers when she first experiences color.

In criticising the hypothesis, Tye says that when Mary has her first experience of red and remarks "so, this is what it is like to experience red", Mary does discover something that the Ability Hypothesis cannot handle. She now knows what it is like to experience red. On the Ability Hypothesis, she has acquired some know-how. But, argues Tye, she retains that know-how, even after she stops having any experience of red. Intuitively "while she attends to her experience, Mary has knowledge-that she didn't have before, knowledge that this is the experience of red". I take it that by 'this' what is meant is 'this

© *A Field Guide to Philosophy of Mind* (Winter 2002)
Symposium on Michael Tye's *Consciousness, Color and Content* : 34-47.
<http://host.uniroma3.it/progetti/kant/field/tyesymp.pdf>

experience'. Tye proceeds to argue that Mary can be thought to acquire, through having an experience, a phenomenal concept, where a phenomenal concept is one that plays a certain distinctive functional role. That functional role, it is further argued, does not provide a threat for physicalism. Tye's account of phenomenal concepts also plays an important role in Tye's treatment of the well-known Explanatory Gap. He argues that the latter is an illusion, one that can be handled, once we have a full understanding of phenomenal concepts.

In my criticisms, I shall concentrate on two chapters, Ch.3 where he describes the representationalist thesis in some detail, and Ch.7, in which the main topic is color. Since I disagree with his thesis, I shall be critical, but I should acknowledge the merits of the book. Tye is to be congratulated for spelling out and defending the theory in great clarity, and with strong arguments.

1.1 Transparency of Experience

One of the powerful motivations for the Representationalist thesis is provided by the transparency of experience. "Focus your attention on the scene before your eyes and on how things look to you. You see various objects by seeing their facing surfaces." (p.45) In seeing these surfaces, he holds, you are immediately and directly aware of a whole host of qualities. You experience these qualities as qualities of the surfaces. You do not experience any of these qualities as qualities of your experience.

So far, this account is non-controversial. What is problematical, however, is what exactly these qualities are, and what exactly is the sense in which one is *aware of* the qualities. Tye has a brief argument to dispose of two hypotheses: (i) that these qualities are qualities of one's (visual) experiences; (ii) that they are qualities of some inner object. The best hypothesis is that they are qualities of external surfaces if they are qualities of anything (and since there are hallucinations and illusions, he holds, they may not be the qualities of anything.)

To suppose that the qualities of which perceivers are directly aware in undergoing ordinary visual experiences, are really qualities of the experiences would be to convict such experiences of massive error, But this, Tye says, is just not credible (it is totally implausible.) Accordingly, the qualities of which you are directly aware, in focussing on the scene before your eyes and how they look, are not qualities of your visual experiences. Moreover, when you introspect your experience, you are not aware of any inner object or thing. "The only object of which you are aware are the external ones making up the scene before your eyes." (p. 47.)

Since you are not directly aware of any qualities of your inner experiences, your experience is transparent to you. But when you introspect you are certainly aware of the phenomenal character of your visual experience. "Via introspection you are directly aware of a range of qualities that you experience as being qualities of surfaces at varying distances away and orientations; and thereby you are said to be aware of the phenomenal character of the experience." (p. 47) By being aware of the external qualities, you are aware of what it is like for you.

What then is the visual phenomenal character? One possible hypothesis, Tye remarks, is that it is a quality of the surface experienced. But that hypothesis is intelligible, he adds, only if it is assumed that the surface is an immaterial one of the sort the sense-datum theorists posited. This hypothesis has already been dismissed. The best hypothesis he suggests, is that visual phenomenal character is representational content of a certain sort: content into which certain external qualities enter.

Two crucial claims emerge from Tye's account:

1. If we stipulate that something is a visual phenomenal quality or quale only if it is directly accessible quality of an experience, then there are no visual phenomenal qualities or qualia.

2. There are qualities of which the subject of visual experiences are directly aware via introspection. They are qualities of external surfaces (and volumes and films) if they are qualities of anything. By being aware of these qualities, we are aware of phenomenal character.

There seems to me to be a number of problems, which I shall discuss in the following sections:

1.2 Two Senses of 'Awareness of'
1.3. Phenomenal Character and Introspective Awareness
1.4 The Block-Searle Objection
1.5 Transparency of Pain
1.6 An Alternative Account
1.7 Projection

1.2 Two Senses of 'Awareness of'

In developing his account of phenomenal character, Tye relies upon an account of introspective awareness of phenomenal character. (I shall examine that account in the following section. Here I want to highlight a potential problem.) Introspection of phenomenal character is said to be a reliable automatic process that takes one from being in one state to being in another, from an experience to a conceptual state. It is a process that takes, as input, awareness of external qualities and yields, as output, awareness that a state is present with a certain phenomenal character.

Unfortunately, in explaining transparency, there seems to be equivocation on "awareness of". In seeing the facing surfaces, the perceiver is said to be "immediately and directly aware of a whole host of qualities as qualities of the surfaces" [p. 46]. Here the awareness consists in having the experiences which are nonconceptual. On the other hand, it is also stated that 'via introspection you are directly aware of a range of qualities that you experience as being qualities of surfaces at varying distances away and orientations". Here the awareness would seem to be conceptual. Introspection is said to be a process that results in your being in a conceptual state: a state of awareness-that.

Now, it may well be that Tye is using the expression "awareness of a quality" ambiguously, once in a nonconceptual sense, once in a conceptual sense, and everything is in order. But it would be better if he didn't. It would be better if he made it clear which sense is being used. This is especially crucial if, as when talking about introspection, he uses the phrase "directly aware of a quality".

I suspect that the ambiguity is not benign. Tye's argument depends on our accepting as a fact a certain phenomenon: the transparency of experience. In asking us to focus our attention on the scene before one's eyes, one in which you see various objects by seeing the facing surfaces, Tye asks us to take as obvious, the fact that one is immediately aware of a whole host of qualities. But which sense of "aware of" is appealed to here?

I am inclined to think that there is a third sense of "awareness of", one that combines elements of the other two. Intuitively, I am presented with an instance of a quality (or complex of qualities: redness, roundness, hard-edgedness, . . . ,) and I am aware of it as, say, being present and before me, and perhaps as an instance of red. In accepting the phenomenon of transparency, I suggest, we are adopting some such sense. I may well be wrong in my account of the intuitive sense, but there seems to be grave possibility of error here in Tye's "intuitively appealing" account.

1.3 Phenomenal Character and Introspective Awareness

In explaining what introspective awareness of phenomenal character amounts to, Tye emphasises that this awareness consists of awareness-that: awareness that an experience with a certain phenomenal character is present (p. 52). Introspection of phenomenal character is said to be a reliable automatic process that takes one from being in one state to being in another, from an experience to a conceptual state. It is a process that takes, as input, awareness of external qualities and yields, as output awareness that a state is present with a certain phenomenal character.

On this account, there are two distinct states: the first is an experience or a feeling or pain, which is an awareness of some objective quality, either a property of an external object (in the case of a visual experience) or a disturbance in one's body, in the case of feelings and pains; the second is a conceptual state: an awareness-that. The first state is a nonconceptual state that is apt to give rise to the conceptual state. If one introspects one's experiences or feelings, one forms a conceptual state: a state of awareness that an experience with a certain phenomenal character is present.

What this second kind of awareness amounts to is that I am aware that there is present something (external) that has certain public objective qualities. These qualities are qualities specific to the representational content of the experiences. It is claimed that I am not aware of any of the intrinsic qualities of this experience itself, I am only aware of the content of the experience, i.e., of the qualities that the experience represents certain objects as having.

One possibility that Tye does not consider seriously enough is that the perceiver might be aware both of intrinsic qualities and content. It may be that the intrinsic qualities

contribute to the content. It could be, for example, as Hume and others have suggested, that we project (or 'project') some of the intrinsic qualities onto the object. Tye does reject the notion of projection, but for reasons I shall discuss later, I think that the notion is defensible.

I have a more serious objection. I think that there is an unresolved tension in Tye's account (a tension shared with the similar account given by Dretske). In what sense is the nonconceptual state the "awareness of an external quality" an awareness at all, an experience at all? It is a state that carries nonconceptual content about the external quality and it is the last stage in a process that results (other things being equal) in a conceptual state, an awareness-that there is something red and moving, say in front of one. But that doesn't make it an experience or awareness. Furthermore, to say that it has a certain intrinsic quality that is causally related in an appropriate way with the external quality. (What exactly the relationship is may be a matter of dispute and it is not easy to specify it in detail: witness the writings on the subject by Fodor.) Is that all it means to say that the state constitutes "awareness of the qualities": that it carries this content? That surely is not awareness in anything like a normal sense. Nor is it remotely like an experience. Why call it awareness or experience then? (I think that John Heil makes a similar criticism of Dretske in his essay in the *Dretske and His Critics* volume.)

Phenomenal concepts are held to be conceptually irreducible concepts that function in the right sort of way. It is part of their characteristic functional role, qua phenomenal concepts, that they enable us to discriminate phenomenal qualities and states, "directly on the basis of introspection". Thus far, so good. But Tye's theory is that the phenomenal qualities that we discriminate are qualities specifiable in the content of the experience, and are not intrinsic qualities of the experience. But if this is so, I find it difficult to understand how there is a difference between what results when I introspect my perceptual experience of a brown shoe on the floor as opposed to what results when I introspect my belief that there is a brown shoe. In either case, all I can discriminate are the qualities contained in the content of the relevant mental state, and on Tye's theory, the content is the same. I know that he believes, as we all do, that the states are different, but once we explain phenomenal character in terms of content, then I cannot see how his theory has the resources to explain the difference. Nor will it help to be told that the content of the perceptual experience is nonconceptual and the content for the belief is conceptual. For it seems to me that when one introspects the perceptual experience, and discriminates the relevant qualities, one becomes aware of the qualities by conceptualising them.

One may be tempted to say that when one has the nonconceptual experience one is aware of a quality, that is, that one is aware of an instance of that quality. But this is not what Tye's theory entails. One is not aware of an instance of the quality, i.e. of an actual instance of that particular quality. One is aware of an experience which has a certain content, that is one is aware that something has that quality. The quality, for Tye, is something abstract.

It may be replied that we can distinguish between two applications of the phenomenal concept. One is a recognitional application: one recognises a current instance of an experience of red as such an instance. the other is a non-recognitional application. But from the point of view of the subject, it is hard to see how the subject can distinguish them. It is hard to see how Tye's theory can explain the capacity that subjects have to distinguish them in the way that they do.

1.4 The Block-Searle Objection:

Tye considers an objection raised by Block and by Searle to the representationalist approach to the phenomenal character of pain. The objection is that the possibility of seeing one's damaged leg, say, while one is feeling pain in the leg shows that the phenomenology of the pain experience cannot be captured by its representational content since the content of the perceptual experience as one views the damaged leg is the same as, or very similar to, the content of the pain experience. The perceptual experience nonconceptually represents features such as color, shape, orientation of surface, presence of an edge, etc. It does not nonconceptually represent tissue damage (or any other comparable quality or cluster of qualities). The leg, says Tye will not look damaged in the phenomenal sense of the term 'looks'.

But how can Tye maintain this? He argues in the case of *looks red*, that given that what it is to be red is to be disposed to reflect certain percentages of light, then if a spot looks red, then the spot looks disposed to reflect such and such percentages of the light. If that is so, why should it not be that the damaged leg look damaged in the phenomenal sense. In any case, it seems to me that the phenomenal sense applies to a range of features or cluster of features besides the ones he cites: there are rusty looks, wooden looks, jarrah looks, metallic looks, and damaged-tissue looks. That is, there is a characteristic experience-type that things which are rusty, wooden, jarrah, metallic, and damaged tissue cause in optimal conditions.

1.5 Transparency of Pain

Tye's Representationalist thesis has considerable plausibility concerning the transparency of, say, visual experience, though I think as Moore suggested, it is possible to give a different interpretation of the transparency. However, I cannot for the life of me see how the experience of pain is transparent. I acknowledge that our experiences of pain have representational content and that they represent a bodily region, say a leg or head or tooth, as being a certain way, and that they *can* be used to represent a bodily disturbance, but how could it be thought that the experience is transparent? The natural account of pain is surely that we are aware of the bodily disturbance through being aware of the intrinsic quality of the pain. The experience of pain has content all right: it represents the intrinsic quality as going on in the body. The ache is experienced in the head, behind the eye, say. Of course, we do not believe that the feeling is there, but the phenomenology of pain is that we experience the pain as if it were there. I feel the pain in the foot just as I feel movement and pressure in the foot. Indeed the feeling of pressure in the foot can quickly become a feeling of pressure and pain in the same region in the foot.

1.6 An Alternative Account:

The example of pain and the possibility of providing an alternative explanation of transparency for visual experience should lead us to take seriously that alternative. That is to say, we should consider the sense-datum alternative (in one sense of 'sense-datum').

Tye considers the sense-datum hypothesis very briefly. The theory is "unacceptable, however, for a whole host of familiar reasons". In my experience most of the familiar reasons usually given are either bad ones or not relevant to the most carefully formulated versions. An example are the reasons given by John Austin in *Sense and Sensibilia*, where he almost exclusively concentrates on A.J.Ayer's version and then often not fairly. (Austin is cited with approval in Putnam's Dewey Lectures as providing the best arguments against sense-datum theories.) I trust that Tye's "familiar reasons" are not these. In my opinion, the arguments presented by Howard Robinson in his *Perception* are very persuasive, if not compelling. They are certainly strong enough to deserve consideration. I do not think that they fail for any of the usually cited "familiar reasons".

It is interesting that at a later stage, (p. 48) Tye asks "what then is visual character?". One possible hypothesis, he writes, is that it is a quality of the surface experienced. That hypothesis, he claims, is intelligible only if it is assumed that the surface is an immaterial one of the sort the sense-datum theorists posited. (By 'intelligible', I assume he means 'plausible or 'reasonable'.) He does not consider further the claim, going on instead to suggest that the best hypothesis (no doubt for avoiding that consequence) is that visual phenomenal character is representational content of a certain sort. I think that he (and we) should consider the alternative more seriously. I would have thought that, intuitively, it is highly plausible that there is such a thing as a visual phenomenal character which is a quality of the surface experienced, and that for the innocent perceiver, it is a quality of an actual physical existing surface. It is captured in the thought that most people have that there is a range of objects (or their surfaces) that have a sensuous phenomenal character: a Perth (or Californian) sunset, a field of poppies, the Pacific Ocean, a film by Visconti. Likewise for other forms of perception: the taste of a ripe peach or of a fine red wine, the sound of a piece of music.

If Tye wishes to reject (as he does) this ordinary intuitive belief, that this character is a quality of an actual existing surface, then it would seem that if it is a quality of anything it is a quality of an immaterial surface. Tye rejects that too, but it is hard to see how the Representationalist account that he espouses can handle the sensuous character of the objects described above.

This point is similar to the point raised in the objection levelled by Block and Searle against the representationalist account of pain. It is particularly significant in handling color experience. It is non-controversial that we experience color as part of the physical world, as on the surface of physical bodies, as spread through volumes and films. One of the central features of these colors is their sensuousness. It is the feature which Tye described in an earlier paper. [Tye (1992), 'Visual Qualia and Visual Content' in Crane

Tim (ed.) (1992), *The Contents of Experience*, Cambridge: CUP, pp. 158-77.] There he considers a hypothetical appeal to first person phenomenology:

Standing on the beach in Santa Barbara a couple of summers ago on a bright sunny day, I found myself transfixed by the intense blue of the Pacific Ocean. Was I not here delighting in the phenomenal aspects of my visual experience? And if I was, doesn't this show that there are visual qualia? [p.160]

He is not convinced. It seems to him

that what I found so pleasing in the above instance, what I was focussing on as it were, were a certain shade and intensity of the color blue. I experienced blue as a property of the ocean not as a property of my experience. My experience itself certainly wasn't blue. Rather it was an experience that represented the ocean as blue. What I was really delighting in then were specific aspects of the content of my experience. It was the content, not anything else, that was immediately accessible to my consciousness and that has aspects I found so pleasing. [p. 160]

But how does the representational account, in his version, handle the sensuousness of the experience? On the face of it, it is part of the phenomenal character of the experience. In the example given, the phenomenal character is part of the representational content, of the experience. But I cannot see how Tye's thesis can handle this situation, given how content is construed. It may be that the content can be specified in terms of the property redness, but we have to think of what redness amounts to, on Tye's theory: it is a reflectance profile. We might point out that this property is not the right kind of property: for one thing it is a dispositional property and if I am enjoying anything it not the dispositional property but the exercise of the disposition. But put that point aside for the moment. There is a more general problem.

It is not clear to me how one is supposed to be aware of this property when I delight in the sensuous character of the colored objects. I am in a certain state, one that carries nonconceptual content about color. What it is for this state to carry nonconceptual content is for it to be of a certain type, one that is causally related (in the right way) with red-type states. This means that on a specific occasion when I delight in the colored experience, the property red may not be instantiated. And even if it is instantiated, how exactly can I be experiencing it?

I am not aware of the intrinsic qualities of the experience, but how can it help that I am aware of the content which is nonconceptual. To be part of the content, a property must be causally related in the right kind of way to my experience, but given Tye's account, to be aware of the property is simply for me to be in a certain state, which stands in a complex relation to instances of the property. It does not count as any normal kind of awareness. Nor does it help to bring in reference to the conceptual state that the nonconceptual state leads to. This second state is a state of awareness-that, but consider what I am aware of: I am aware that I am having an experience which represents an object as being a certain way. But it is hard to see how this kind of awareness is anything I could rejoice in. Why should it make any difference to me that I

become aware that the sunset is red, any more that I learn from reading it in the paper that the sunset at 5 p.m. is red. Of course they have different causes, but given that in each case the conceptual state is an awareness-that, there does not seem to be any significant difference between the two.

The point might be put this way. The nonconceptual content of an experience is informational content. It must be more that information, but it is at least that. What makes it representational is related to the causal history of the state. The quality represented is one that has had a role in the person's acquiring that kind of state, and the use that the perceiver can make of that information. But what can that quality have to do with my current experience. It seems to me that what is required is either a current instance of the quality is there for me to be aware of, or alternatively the perceiver is disposed to act towards the object bearing the quality in a way that is sensitive to its presence.

1.7 Projection

In arguing that the phenomenal character of an experience is given by the content of the experience, Tye does not consider that the perceiver might be aware both of intrinsic qualities and content. It may be that the intrinsic qualities contribute to the content. It could be, for example, as Hume and others have suggested, that we project (or 'project') some of the intrinsic qualities onto the object. Near the end (p. 165), Tye says that projectivism, upon reflection, seems incomprehensible. He claims that he does not understand how subjective qualities can be projected, and especially onto physical objects. The answer to this is that the projection is not literal, but metaphorical: it is a case of 'projection'. One example is that a picture of a man in a red coat can be such that the red in the picture is used to represent the red of the man's coat. That is to say, the red in the picture exemplifies the red of the coat represented. Likewise two actors kissing on stage can exemplify two characters kissing. In the case of pain, the pain one feels in one's leg is a subjective feeling that one 'projects' onto the leg. It is not a real projection. One has a body image which represents the body. The pain is projected on to, is located on, that part of the body image that represents the leg. Likewise with colors. There is a subjective quality which one 'projects' onto an external object, say to the moon, to represent it as yellow. One does not actually projects the quality onto the physical moon: one projects it onto that part of one's subjective visual field that represents the moon. The basis of the metaphor is that the perceiver automatically and naturally takes the representation of the leg, to be identical to 'that which I reach for when trying to ease the pain', i.e. the part of the leg that is the source of the pain. Likewise, s/he takes the moon-representation to be identical to 'that which I point at when indicating the moon'.

Indeed, the notion of projection seems perfectly fitted to Tye's Representationalism, or to a suitably modified version of it. The intrinsic quality of the subjective experience need not be thought of as projected onto a physical object, or onto the subject's leg: it is 'projected' onto the representational content of the experience. That is to say, it contributes to, or is a part of, the content, in the way that a property of a photograph *may* contribute to the content of the photograph.

2. Color

2.1 Tye's position on Color

In the chapter on color, Tye begins by describing a view that has a distinguished history, and has been recently defended by a number of different theorists: Cosmides and Tooby, Boghossian and Velleman, and L. Hardin. The view is that science has rebutted the common-sense view of color. This commonsense view is that colors are objective, perceiver-independent qualities of physical objects.

Tye argues that the scientific facts allow for another proposal consistent with the commonsense view: colors are objective, microstructural properties, e.g., reflectances. It is important that such a proposal be defended since, as Russell pointed out, and Tye agrees, if things in the world lack color, then the only way to avoid the conclusion that there is a deep and pervasive error in our color experiences is to claim that things we experience are not outside the mind at all.

Tye's treatment in this chapter is largely confined to treating some of the arguments developed by some of the theorists cited above. I do not propose to discuss these arguments apart from the one by Hardin because I think that there are other powerful reasons for thinking the commonsense view false. Tye's discussion is relevant, however, for I think that in this discussion he makes a number of questionable assumptions.

There seem to be two major elements in his characterisation of the commonsense conception:

1. The colors we see objects and surfaces to have are observer-independent properties of those objects and surfaces. We think of colors as inhering in surfaces and in volumes and films. We take it for granted that objects typically retain their colors when not seen, thereby helping us to re-identify them,

2. Another important fact about color which is manifest to us in our everyday life is color constancy. Objects do not typically appear to change their colors during the day as the sunlight changes.

There are, Tye thinks, three objective accounts consistent with this conception, but there are other reasons for ruling out two of them. The most plausible account, he holds, is Reductive Physicalism: the thesis that colors are physical properties whose natures are discoverable by empirical investigation. The version of this broad thesis that he seems to favour is some form of Anthopocentric Objectivism. (Colors are identified with some objective properties which in themselves are of little interest except in so far as they have a distinctive effect on those organisms, especially humans, built to respond to them.) In the case of surface colors, reflectances are the natural candidates.

My major criticism of Tye's approach is that the conception of the commonsense view of colors that he is working with, is under-described. It is clear that the two elements Tye attributes to the common-sense conception are not enough, they do not enable us

to distinguish the concept of color from many other concepts. There are, that is to say, important elements other than the two he describes. Once we take these elements into consideration, it will become apparent that we will have trouble retaining the commonsense conception. Although he does consider some of them in his discussion I do not think that he does enough to justify dismissing them. I think that some of these elements were in the background of the thinking of many of the traditional philosophers and scientists, and if they were not, they should have been. They should be present in our thinking.

2.2 Reductive Physicalism

In order to defend reductive Physicalism about colors we need to give an accurate characterisation of the common-sense conception. In order to do that we need to say something about how the ordinary concept of color operates, e.g., about the practices, (discursive and non-discursive) in which the concept is embedded.

We should note that in many scientific microstructural reductions, what happens, in effect, is that the scientists, or the philosophers who do the reducing on their behalf, redefine the prior concept and having replaced it with a new concept, then identify the property picked out with the new concept with a new microstructural property. Of course it is not arbitrary that this is done, but given that it is done, it is clear that it is being assumed that the prior concept is inadequate and needs to be replaced. Tye can hardly apply this approach to reductive Physicalism for colors, since it involves rejecting the commonsense conception of colors. He wants to argue that reductive Physicalism is consistent with that conception. (As Searle has pointed out, the reconstruction of the prior concept often omits elements that have to do with subjective elements. To omit these in the case of color hands the reductionist a pyrrhic victory.)

If physical reduction for colors is to be consistent with the common-sense conception, then it needs to respect the way the concept of color operates. That is, it must acknowledge that the concept operates within a certain set of conceptual practices, central to which are the naming practices, learning and teaching practices. If color terms are taken as functioning as names for certain properties, then there are different ways in which they might do so. There are three primary candidates we have to choose from: causal reference-fixing models, descriptivist models and models which combine descriptive and causal elements. These theories offer different models for understanding what kind of properties color terms name, and how they name them. Putting the point in relation to concepts, we can say that the color concepts expressed with terms such as “brown”, “olive”, “turquoise”, “crimson”, etc., conceptualise colors as certain kinds or properties. The different models offer different ways for the conceptualisation to work.

It seems to me that since there are different models that might apply, the reductionist must defend their choice. Tye does not do so but rather presupposes one of the models. More to the point, I think that there is sufficient reason to think that he has made the wrong choice.

Problems:

There are several problems. In the first place Tye assumes a certain view about how color terms operate, e.g. in terms of a certain way of construing the sense-reference distinction. That is controversial. It seems to me that our ordinary color terms function in such a way that they are understood to be ostensively defined. One teaches and learns the terms through ostensive situations and through dealing with examples of colors.

If we accept this view, we need to say something about how the ostensive definition should be understood, or at least filled out. that is to say, we need to characterise the ostensive definition I such a way as to be faithful to the teaching and learning practices. It seems to me that the more faithful way is to characterise red, say, as "this manifest feature exemplified here (and in blood, and sunsets, and John Bull coats, and Essendon football jumpers, and so on", and not as "red is the hidden feature responsible for the way that blood, sunsets, . . . appear".

I do not see the problem with saying that our ordinary practices assume that color is a manifest, rather than a hidden, feature. If color is not a manifest property, what is? And surely the fact that we theorists function very well with a notion of hidden essences (i.e., with properties whose natures are non-manifest) shows that we have a viable concept of *being manifest*. Given the centrality of the ostensive teaching and learning practices, we can identify other features essential to the characterisation of the commonsense conception: colors are intrinsic, non-relational and non-dispositional features. Finally, as was argued in section 1.6 above, there is the sensuous character of the objective color properties that must be acknowledged.

There is a second important problem for Tye's account. The microstructural properties associated with the various colours are very different for surfaces than they are for lights, and as they are for volumes or for films, or for spots of light on the wall, or for objects such peacock feathers where diffraction effects operate. For the ordinary concept, it does not matter much whether we are talking about houses, or bird feathers or sunsets, or skies or glasses of wine, we use the terms blue, red, burgundy, claret etc to apply to all of these. That we do so is compatible with our recognising that there are certain important differences in the ways these colors appear. There is no question that sky blue is similar (enough) to a blue-bird's coat or a Cambridge rugby jumper. This point is important. In teaching situations we do not hesitate to use examples for all of these kinds of objects, to teach 'blue' for example with respect to the sky, the ocean, the blue in a peacock's tail, the blue on a flag, and so on. If we mix lights and cast them on a screen, there is little hesitation in identifying them as yellow, blue, orange, and so on.

As I pointed out earlier, it is clear that the two elements Tye attributes to the common-sense conception are not enough, that they do not enable us to distinguish the concept of color from many other concepts. Once we have filled out the characterisation of the commonsense conception in more detail, we can see that there are crucial features that Tye has ignored. There are two of particular importance. The first is that colors are treated as properties whose natures are manifest. The second is that volumes, films and

surfaces share the same property. The first element rules out the possibility that according to the commonsense conception the natures of colors could be hidden. The second rules out the possibility that the nature of surface color could be different for that of volume color and film color etc. This fact precludes (identifying or reducing) the property picked out by the commonsense concept with any microstructural property. The objectivist it seems to me is forced to reject the common-sense conception.

There are other features that are also important: that colors are intrinsic, non-relational, non-dispositional and sensuous properties. Even without them, but especially with them, the commonsense concept is in trouble.

2.4 Unitary/Binary Structure of Colors

Tye considers and rejects Hardin's objection to physicalism with respect to colors: some colors are unitary, in particular red, yellow, blue, black and white. Others are binary, for example orange, pink, purple and blue. Tye argues that, contrary to what Hardin contends, this provides no problem for physicalism. To handle it, however, Tye treats orange, for example, not simply in terms of having a certain reflectance profile, i.e., as being disposed to reflect different proportions of light from the range of wavelengths associated with standard illumination. Rather it is reflectance which has a certain power defined in terms of the effect it has on the perceiver's opponent-processing channels (suitably specified, as in the best available theory). Tye may well be right that in this way a theorist can explain the unitary/binary structure of color appearance. Hardin's point surely was directed against certain objectivist conceptions of color: those that construe colors as objective, perceiver-independent properties of objects, that is, as either intrinsic properties of surfaces (volumes, etc.,) or as dispositional properties such as reflectances. These properties could not explain the underlying unitary/binary structure. It might be possible to reconstruct a new physicalist property that was no longer perceiver-independent, i.e, a property defined in terms of its effect of the perceiver's opponent-processing channels. Perhaps Hardin's argument fails against that view, (without further argument), but that view was not the original target. More importantly for Tye's purposes, however, such a concept would be different from the commonsense one.

Moreover, if we recall Tye's characterisation of the transparency of experience, it is hard to see how this concept of color fits his description of the transparency of experience. There he argued that in seeing the facing surfaces of objects before you "you are immediately and directly aware of a whole host of qualities. You experience these qualities as qualities of the surface. You do not experience any of these qualities as qualities of your experiences." [p.46] If we can be so confident of what we experience here, that the qualities are not qualities of our experiences, for example, is it all plausible that these qualities that we claim to experience are the reflectance properties of the type that Tye requires. I suggest not.

2.5. Conclusion

I do not think that Tye has done enough to justify his claim that his version of Representationalism can explain our experience of color, and in particular the phenomenal character of color. It seems to me that his account of the commonsense conception of colour is underdescribed. Once we fill out that conception in an adequate manner, I claim, the difficulties for Tye's thesis clearly emerge.

Phenomenal Character and Color - Reply to Maund

by Michael Tye

Department of Philosophy
The University of Texas at Austin
Austin, Texas 78750
USA

1) Barry Maund covers a lot of ground in his essay, and he says much of interest. Let me begin with some of his remarks about transparency. According to Maund, in my discussion of transparency, there is an equivocation in my use of the expression "awareness of". I would say rather that in the following sentence, partly quoted by Maund from *Consciousness, Color, & Content*, I misplaced the qualifier 'via introspection':

Via introspection, you are directly aware of a range of qualities which you experience as being qualities of surfaces at varying distances away and orientations and thereby you are aware of the phenomenal character of your experience.

I should have placed 'via introspection' after the words 'and thereby'. The first part of the remark pertains to awareness in seeing the surfaces before the eyes, as was made explicit in my comments in the prior two pages. The first 'aware of' in the passage expresses nonconceptual awareness. The second 'aware of' expresses awareness-that. The abstract noun 'the phenomenal character of your experience' stands in for a factive clause: you are aware that your experience has so-and-so phenomenal character. This is true for abstract nouns generally when they follow verbs for awareness of any sort. Consider, for example, 'the sentence the man spoke' in the case that I am said to hear the sentence the man spoke. It does not suffice for the remark about me to be true merely that I hear the sentence. I must hear *that* the sentence the man spoke was such-and-such.

The general point here is that if you try to introspect a visual experience, say, you will certainly become aware *that* you are directly aware of various surfaces and qualities, which you experience as being qualities of the surfaces, but you will not find yourself being aware, as you introspect, of an inner token experience or of any qualities of an inner experience. By being aware of the qualities and surfaces outside (or apparently outside), you are aware that you have an experience with a certain phenomenal character. That is all.

2) Maund moves next to what he calls "a more serious objection." He asks rhetorically

© *A Field Guide to Philosophy of Mind* (Winter 2002)
Symposium on Michael Tye's *Consciousness, Color and Content* : 48-53.
<http://host.uniroma3.it/progetti/kant/field/tyesymp.pdf>

whether all there is to a state's being an experience is that it carry nonconceptual content and that it be "the last stage in a process that results (other things being equal) in a conceptual state." A couple of sentences later, he asks again rhetorically whether this is "all it means to say that the state constitutes "awareness of the (relevant) qualities".

On my theory, what it is for a state to be an experience is that it be a PANIC state. That is indeed all there is to a state's being an experience. However, nowhere do I claim that the meaning of the term 'experience' can be cashed out in terms of PANIC. Nowhere do I claim that the concept EXPERIENCE is the same as the concept PANIC. Indeed, I explicitly say the contrary, and I give my reasons (see *CCC*, pp. 53-4). Phenomenal concepts, I maintain, are conceptually irreducible; and the concept EXPERIENCE is one such concept of a very general sort. So, Maund's second question is not to the point.

3) Maund wonders how, on my account, "there is a difference between what results when I introspect my perceptual experience of a brown shoe on the floor as opposed to what results when I introspect my belief that there is a brown shoe." He continues:

In either case, all I can discriminate are the qualities contained in the content of the relevant mental state, and on Tye's theory, the content is the same.

I do not know why Maund says this. The thesis of transparency is a thesis held with respect to experiences. It is not a thesis held with respect to beliefs. When I introspect a perceptual experience, I am directly aware of various qualities represented by the experience, for in undergoing the experience I am directly conscious of the qualities. And what introspection tells me is simply *that* I am undergoing the relevant perceptual experience. Introspection is not awareness of the experience. Nor is it awareness of qualities of the experience.

When I introspect a belief, or better its conscious manifestation in an occurrent thought, I am not directly aware of the qualities represented by the belief. I am simply aware that I have the relevant belief. Moreover, the content of the belief is certainly not the same as the content of the experience in the first case. The latter is nonconceptual and very rich; the former conceptual and nothing like as rich. Nor is the content of the higher-level awareness in the two cases the same. The awareness in the first case exercises phenomenal concepts. The awareness in the second does not. Further, the second awareness uses the concept of belief; the first does not.

4) Maund approves of the objection raised by Block and Searle that my view cannot distinguish between the phenomenology of seeing one's damaged leg and feeling pain there. I think otherwise. One's visual experience, as one views the leg, nonconceptually represents such features as color, shape, orientation of surface, presence of an edge. It does not nonconceptually represent tissue damage. One's pain does nonconceptually represent tissue damage, but it does not represent the other features.

I concede that, going by the phenomenal look of the leg, one will judge it to be damaged, and thus that it will look to be damaged. This is a conceptual use of the term 'look', however. In general, if X looks to be F to person P, P must possess the concept F. The phenomenal sense of 'looks' is nonconceptual. It is captured by 'looks F', where 'F' is a term for a quality of which one is directly aware as one undergoes the relevant experience.

Maund objects that the damaged leg does look damaged in the phenomenal sense of 'look'. He comments:

....it seems to me that the phenomenal sense applies to a range of features or cluster of features besides the ones Tye cites: there are rusty looks, wooden looks, jarrah looks, metallic looks, and damaged-tissue looks.

I disagree. If one views something rusty in standard conditions, one is directly aware of a range of color and texture qualities, on the basis of which one judges that it is rusty. The object looks to one to be rusty; moreover it looks like other rusty things, and indeed, in one ordinary way of speaking, it looks rusty, but it doesn't look rusty in the nonconceptual, phenomenal sense. For rustiness isn't a quality of which one is directly aware when one introspects one's experience any more than is the quality of being feline, when something looks feline to one. Intuitively, felines and twin felines (molecule by molecule duplicates of cats belonging to a different species on a variant of Putnam's planet, twin earth) are phenomenally indistinguishable. In the phenomenal sense, they look alike. But cats look feline to us, whereas twin cats look twin-feline to the inhabitants of the other planet.

What is true here for cats is true mutatis mutandis for rusty things, wooden things, etc. And it is true for damaged tissue as well. Imagine in this case that on the twin planet, there is no tissue but an artificial look alike. Note, incidentally, that this is not to deny that "there is a characteristic experience-type that things which are rusty, wooden, jarrah, metallic and damaged tissue cause in optimal conditions" (Maund). There is such a type; but at the nonconceptual phenomenal level, that type is individuated by the cluster of qualities of which the subject is directly aware via introspection (qualities that are also represented by the experience). And those qualities are at the level of shape, color, texture etc. They do not include rustiness, woodenness, and so on.

5) Maund says that he finds the transparency view very implausible for pain. He claims that the natural account of pain is that "we are aware of the bodily disturbance through being aware of the intrinsic quality of the pain." A little later he remarks:

....the pain one feels in one's leg is a subjective feeling that one 'projects' onto the leg. It is not a real projection. One has a body image which represents the body. The pain is projected on to, is located on, that part of the body image which represents the leg. Likewise with colors. There is a subjective quality which one 'projects' onto an external object, say to the moon, to represent it as yellow.

I confess that do not understand any of this. If I see the moon, I am not aware of a subjective visual field that represents the moon. I am aware of the moon and perhaps some stars located in distant regions of space before my eyes. Likewise, if I have a pain in my leg, I am not aware of an image that represents my leg. I'm aware of my leg and its condition. To suppose that it is the representation itself -- the subjective visual field or the body image -- of which I am really (directly) aware in these cases is like supposing that if I desire eternal life, what I really (directly) desire is the idea of eternal life. That, however, is not what I desire. The idea of eternal life I already have. What I desire is the real thing. And it does not help, of course, to say that it must be the representation of which I am aware, since the case might be one of hallucination — no moon or no leg — for patently, if there is no eternal life, it still isn't the idea of such a life that I really desire. If the pain is a phantom one or the visual experience totally delusive, I simply undergo an experience which represents something that isn't there.

It seems to me, then, that the right thing to say is that when I attend to a pain in my finger, I am directly aware of a certain quality or qualities as instantiated in my finger. Moreover, and relatedly, the only particulars of which I am then aware are my finger and things going on in it (for example, its bleeding). My awareness is of my finger and how it feels. The qualities I experience as bad or unpleasant are ones the finger or part of the finger or a temporary condition within the finger apparently have. My experience of pain is thus transparent to me (or so I continue to hold). When I try to focus upon it, I 'see' right through it, as it were, to the entities it represents.

6) Maund insists that visual phenomenal character is a quality of objects — not a real property but an experienced property, one that visual experience projects on to objects. He says that it is a property of the Perth sunset and a field poppies, for example.

I agree that the qualities I find so pleasing in a sunset or a field of flowers are ones I experience the sunset of flowers as having. But why insist that these are qualities the experiences project upon the objects? If a particular poppy is vivid red, why not say that it is the vivid redness of the poppy in which I delight? By experiencing a quality I like so much, I undergo an experience whose phenomenal character has an aspect, or component part, that is delightful to me. For the phenomenal character, on my view, is a certain representational content my experience has and that content is one into which the vivid redness of the poppy enters.

Maund objects that I don't delight in my awareness that I am having an experience that represents the object a certain way. I concur. For one thing, as just noted, it is an aspect of the content of the experience that pleases me. The experience is the bearer of phenomenal character, not the higher-level awareness that the experience is present. For another, it isn't any aspect of the content of the experience, *qua* aspect of the content, in which I take delight. Again, that phenomenal character is a certain sort of representational content isn't something given to me directly in experience itself. As I note in the book, representationalism is best viewed as a hypothesis that is justified in terms of its explanatory power.

7) The last part of Maund's essay concerns my view of color. According to Maund, colors are manifest qualities, not qualities with a hidden nature. I agree with Maund that colors are manifest. After all, we see things by seeing their facing surfaces and we see the facing surfaces by seeing their colors. In this way, colors contrast with such qualities as being an electron or being a quark. But it does not follow from this that colors do not have a hidden nature. The thesis of revelation — that the nature of color is wholly given to us in sense experience — is much stronger than the thesis that colors are manifest qualities. Revelation, it seems to me, is a philosophical thesis that is no part of common sense. We should all agree (obviously) that colors are not given to us in color experience as having a hidden nature. But we need not also agree that colors are given to us as not having such a nature.

8) Maund says that volumes, films, and surfaces, in being the same color, share the same property. An obvious addition to this list is light sources. Maund takes it for granted that the reflectance view of color is in serious difficulty here; but he is, I think, too hasty. Byrne and Hilbert , in a forthcoming essay, ("Color Realism and Color Science", *Behavioral and Brain Sciences*), propose that reflectance be characterized not in terms of the light *reflected* by a surface but by the light *leaving* it (by reflection, transmission, or emission). Saying that the light leaving a surface is the light the surface *produces*, they identify the reflectance of a surface with its disposition to produce a certain proportion of the incident light. This characterization of reflectance is equivalent to the one I give for opaque, non-luminous surfaces; but for surfaces that emit or transmit light, it gives very different results. As they show, the colors of volumes and light sources can now be accommodated in a relatively straightforward way.

9) On my view of color, the colors we see are entirely objective properties, though they are anthropocentric. Maund takes me to hold that colors are reflectances with powers "defined in terms of the effect [they] have on the perceiver's opponent-processing channels." So, he claims, on my account, colors are not objective, perceiver-independent properties. And this allegedly shows that I have not answered Hardin's point that the objectivist about color cannot preserve the unitary/binary distinction, contrary to my claims in the book.

Maund here has failed to come to grips with my theory. Colors, in my view, are just as perceiver-independent as shapes. The surface colors we see are spectral reflectances that dispose their possessors to produce (via reflection) certain percentages of long, medium, and short wavelength light in certain objectively specifiable viewing conditions. These objective properties of surfaces are ones to which our visual systems are 'tuned'. We find them special because of how our visual systems are constructed, but their being so does not make them subjective. They are real, external, objective properties even though they are of no interest to creatures lacking our visual systems.

My proposal is about the nature of color. It is not, as Maund suggests, a proposal about the concept of color. I am certainly not offering definitions of color words or analyses of color concepts in terms of the effects certain properties have on the opponent processing systems of perceivers. Nor am I offering an account of the nature of colors that involves perceivers in that nature.

10) Another of Maund's objections is that the view that colors are reflectances doesn't fit with the thesis that color experience is transparent. For, he claims, it just isn't plausible that any of the qualities of which we are directly aware in seeing the facing surface of an object are reflectances. But why not? To be sure, we are not aware of any of these qualities as reflectances. Indeed, on my view, colors are not presented to us in sensory experience under any mode of presentation at all. Our awareness is direct. Maund offers no further argument against this view.

Some Awkwardness in Poised Content?

by William Seager

Division of Humanities
University of Toronto at Scarborough
Scarborough, Ontario
M1C 1A4 Canada

Although the problem of consciousness seemed to be something of a side-issue in the development of naturalistic accounts of mind in the mid to late twentieth century, real progress has been made recently. We now possess, thanks to Michael Tye and other philosophers, an impressive variety of sophisticated theories of consciousness. To my mind, the most significant of these are the representational theories of consciousness.

Michael Tye's *Consciousness, Color and Content* (2000) expands, extends and defends the theory of consciousness which he advanced in *Ten Problems of Consciousness* (1995). Tye's theory, as a representational theory of consciousness, asserts that what consciousness *is* is the active, or at least potentially active, presence within a cognitive system of a set of representations with certain well-defined properties. This is, in one way or another, common to all representational theories. But for Tye, *all* of consciousness is a matter of representational content. Thus even qualitative states of consciousness, the taste of coffee, smell of rose, twinge of pain, etc. are to be understood as representational content, though of a rather special sort. This feature of Tye's view is controversial even within the representational camp, since the 'phenomenality' or qualitative nature of mental states can be regarded as a property of such states in its own right rather than a matter of what or how non-phenomenal features are represented. The prospect of explaining the qualitative aspect of consciousness is the most exciting feature of Tye's representational approach.

Of course, not everyone agrees that the best approach to consciousness is so thoroughly representational (Ned Block has presented several arguments against the representational view; see Block (1990) for example), but I do not want to question the representational theory's basic premise here. Instead, I will raise some issues from within the representational camp.

There are now several versions of the representational theory of consciousness. An important distinction within the field is between 'higher-order' and 'first-order' theories. The former assert that a state, S, is a conscious state if it is the object of a higher-order mental state - a state which represents or is about the original state. Important distincti-

Symposium on Michael Tye's *Consciousness, Color and Content* : 54-64.
<http://host.uniroma3.it/progetti/kant/field/tyesymp.pdf>

ons can be drawn within the higher-order theories as well. There are higher-order *thought* theories, in which consciousness is the result of a higher order state which is a thought about the lower order state (see Rosenthal 1986); other theories prefer to regard the higher order state as a kind of *perception* of the lower order state (see Lycan 1996).

It will be important for us to also distinguish between *actualist* and *dispositionalist* versions of these theories depending on whether a state's being conscious depends on the actual presence of the requisite higher order state or whether a mere disposition to bring about the higher order state is sufficient for the lower order state to be a conscious state (a very well worked out dispositional higher order thought theory of consciousness has been developed by Carruthers (2000)).

First order representational theories of consciousness, of which Tye's theory is an example (see also Dretske 1995), do not require any relation to a higher order state which is about a lower order state for that state to be a conscious state. Since all of the theories at issue here tend to accept pretty standard cognitive science accounts of the mind, there is a huge supply of representational states presumed to be active within any working cognitive system. While higher-order theories have a clear account of what makes a particular representational state conscious, they face a problem about distinguishing mere consciousness from introspective consciousness. On the other hand, first order theories face the non trivial task of distinguishing those representational states which are states of consciousness from the plethora of representational states that are not. This leads, in Tye's theory, to a constraint upon which representational states are conscious somewhat similar to that imposed within higher-order theories.

However, before proceeding with this line of thought, which leads deep into some core issues of representational accounts such as Tye's, I want to digress briefly and introduce a small, and what I hope Tye would regard as a friendly, amendment to his account of conscious pain. It does not strike me as correct that animals cannot 'suffer' pain (see p. 182). This could be regarded as a merely verbal matter if Tye wishes to *define* suffering as being the introspective awareness of pain. But I see nothing to recommend this linguistic legislation. Normally we regard the suffering as what is *bad* about being in pain, and if animals cannot suffer then there doesn't seem to be anything bad about their being in pain and their pain loses (at least a major part) of its moral significance. It also seems intuitively compelling, to me at least, that there are pains which are so intense that they can destroy one's ability to introspect without destroying one's consciousness and while preserving the suffering. These pains are themselves so vicious that they obliterate any awareness of them as states of mind, leaving only the searing awfulness - which is certainly a state of consciousness - which they intrinsically possess.

Of course, on a first order representational account of consciousness whatever qualitative feature states of consciousness possess 'intrinsically' is a matter of what they represent. Thus I suggest that the theory recognise the representation of 'evaluative properties' as features of states of consciousness. Evaluative properties are properties of things (that is, intentional objects of states of consciousness), or putative properties of

things, which make them *good* or *bad*. A pool of water, to a thirsty man, is represented as clear, cool, wet and *good-to-drink*. There is no *inference* to this 'primitive attractiveness' of water when thirsty, though there might be an inference that this water is, in truth, *not* good to drink (because, say, it is suspected to be poisoned). A pain is a representation of a certain part of the body being damaged in a certain way, but also possesses a distinctive negative evaluative representational component which is its painfulness. Suffering can then be defined as being in a state of consciousness with that kind of negative evaluative representational component.

Evaluative properties are complex, since they have a built in relativity to the kind of being representing them and the state of the being doing the representing. Water does not always look good to drink, but when it does there is a distinctive element of consciousness which cannot be neglected. To take another and well known example, consider the analgesic effect of opiates such as demerol. It is common for these to relieve pain, as we say, without actually obliterating the consciousness of the state of the body which has been damaged. Jeffrey Foss provides a harrowing autobiographical account and enlightening discussion of the biochemistry of the curious analgesic effect of demerol on *serious* pain in Foss (2000). If the representation of evaluative properties is included in our theory of consciousness, this effect of demerol has a straightforward explanation. Crudely put, demerol works by blocking the representation of the distinctive negatively evaluative property central to pain. After the injection of demerol, one's cognitive machinery is still representing the damage to one's body; one still can introspect the pain, and know (directly) that it is pain, but it has changed, and the change is that the pain is not so *bad* anymore. Foss describes the phenomenology of this as 'my leg still screamed, but I was no longer inclined to pay any attention ... my intellectual comprehension just did not translate into caring, much less action' (p. 146). In fact, one might speculate that the consciousness of this kind of value is the most basic feature of sensory consciousness, the feature that drove the refinement of the more familiar sensory modalities as a way to facilitate the indirect, behavioural, alteration of these fundamental evaluative aspects of consciousness (in a slogan: approach the good, avoid the bad).

I will not try to develop this idea any further here, though I think it is central for the extension of the representational account to more complex forms of consciousness. I do want to note some of the areas in which it is important. It is crucial for the proper treatment of the consciousness of various kinds of value, and, not unrelated, for the correct account of emotional consciousness. It is also important for the extension of the theory's account of introspection beyond that of purely sensory states (see Seager 2000). Furthermore, it is I think necessary for the proper account of the difference between states of consciousness which are *motivating* and those which are not; crudely, it is the representation of evaluative properties which underlies the way perception, and other forms of consciousness, motivate us to actually do something rather than passively observe. Finally, if somewhat speculatively, the acceptance of evaluative properties as features of states of consciousness suggests an interesting and naturalistic approach to ethics or at least 'moral consciousness'. However, the main point I want to make here is simply that the incorporation of evaluative properties is entirely within the spirit of the representational theory of consciousness and relieves representationalists of

the unpleasant and implausible proposition that animals, and children, cannot suffer even when in obviously excruciating pain.

Now, to return to the main line of argument, recall that I asserted that Tye's theory has a feature surprisingly similar to that of higher-order theories, namely a requirement for a state's being conscious that it bear a certain relation to *thoughts*. To explain this, note first that Tye's theory appears to be restricted to an account of sensory consciousness, conceived sufficiently broadly to allow bodily sensations to count as sensory states. So, if conscious *thinking* involves a kind a consciousness which is non-sensory, which intuitively appears to be the case, Tye's theory does not apply to conscious thinking. Thus there is something of a problem generating a unified account of consciousness from Tye's perspective. This difficulty does not arise for higher order theories It is worth noting, however, that another problem of unification afflicts at least some higher-order thought theories. Some of these (see Rosenthal (1986) for example) regard phenomenality as a feature of mental states in its own right, one that can characterize mental states whether or not they are conscious. Thus there seems to be a kind of disunity in their treatment of phenomenal consciousness versus conscious thought.

In any case, Tye asserts that for a sensory state to be conscious its representational content must be available to more complex cognitive mental states which mediate between the sensory content and behaviour. We might call such states 'higher level' to distinguish them from 'higher order' states which have lower order mental states as their intentional objects. Beliefs and desires would be the most natural examples of such higher level mental states but Tye is clear that cognitive states that perhaps do not fully merit the status of beliefs and desires will suffice to underwrite states of sensory consciousness. Thus certain animals which we may doubt have full-fledged beliefs and desires can still enjoy states of sensory consciousness insofar as they use the content of their sensory representations to facilitate at least somewhat intelligent, unfixed, non-automatic and learned behaviour. Tye goes so far as to include, for example, honeybees in this category.

Obviously, if one makes the demand that sensory states interact with higher level states in order to be conscious, then one can espouse either an actualist or dispositionalist version of the demand. Tye opts for the dispositionalist account. One of the conditions of the PANIC theory is that the content of sensory states must be 'poised' or available to the higher level cognitive states (recall that the other conditions are that the content must be abstract, non-conceptual and intentional). Tye says that such states 'stand ready and available to make a direct difference to beliefs and desires' (p. 172). In Tye (1995) the discussion of *poised content* is even more explicitly dispositionalist; there Tye says that 'to claim that the contents ... must be *poised* is to be understood as requiring that these contents ... stand ready and in position to make a direct impact on the belief/desire system. To say that the contents stand ready in this way is not to say that they always do have such an impact' (1995, p. 138). Poised content need not actually create or modify any beliefs or desires in order to be conscious.

Dispositionalist accounts of the conditions of consciousness face an immediate and I think serious, if rather abstract and 'purely philosophical', objection. Take a subject, S1,

and consider the set of PANIC states of S which are conscious but which in fact have no effect on any higher level cognition. While these states have all sorts of effects on S1 and S1's behaviour they are conscious solely in virtue of their unexercised dispositions to affect high level cognition. Now, take a second subject, S2, who is identical to S1 save that S2 has been modified by the attaching of a device that would block the relevant disposition - that is, make the content unavailable to higher level cognition, but *only* for those states which in fact are *not* going to affect higher level cognition. Note that this device will never have to do anything in S2 (it will only operate, so to speak, in counterfactual situations). There will be absolutely no difference in the neural processes of S1 and S2, nor in their behaviour. The only difference is that S2's brain has within it a totally inert disposition blocking device. Nonetheless, Tye's theory asserts that S2 will have quite different states of consciousness compared to S1. This seems to me extremely implausible. S1 and S2 will behave exactly the same way and they will have exactly the same neural processes at work within their brains. They will even have exactly the same representational states active within them and active in exactly the same way in each of them. Yet S1 supposedly has many more states of consciousness than S2.

It might be replied that S2's brain is 'abnormal' because of the attached device. This kind of abnormality seems irrelevant. Imagine a similar device affixed to S1's brain which is missing its battery, so it cannot operate. S1's brain is now 'abnormal' in the same was as S2's, since after all the device in S2's brain will as a matter of fact never operate. Surely attaching such an inoperable device to S1 will not alter S1's consciousness in any way at all. Why should attaching an almost exactly similar device save that it *could but in fact will not* work make such a huge difference in consciousness?

This objection stems from and supports the intuitively attractive idea that consciousness is an occurrent phenomenon which depends only upon the current state of the subject, and of course there have been attacks on this intuition (for example, Dretske's famous or infamous denial of consciousness to Swampman). To my mind, the intuition seems on a sounder footing than the theories that deny it.

In any case, the dispositional aspect of poised content must be distinguished from introspectibility. There are, to be sure, examples of people who don't 'notice' their own states of consciousness. But this point is a point about the difference between consciousness and introspection. First order theories have an easier time explaining the difference between mere consciousness and introspection than do the higher order theories, since the latter explicitly invoke a higher order mental state which is about the lower order, conscious, state and it is tempting to equate such a higher order state with introspective access to the conscious state. Thus higher order theories have to impose a distinction upon the set of higher order states which divides them into the merely 'consciousness creating' and 'introspective access providing' states. First order theories can, so to speak, borrow the machinery of higher order theories as the basis of their theory of introspection without modifying their theory of consciousness itself.

Thus, introspection is the having of thoughts about one's states of consciousness. But mere consciousness itself requires only the sort of interaction with high level cognition

adumbrated above. One example much discussed by all these theories is that of the distracted driver. This is the phenomenon, of which the reader very likely has first hand experience, of suddenly discovering that one has seemingly not been paying any attention to driving one's car for a disturbingly long period of time. This phenomenon is subject to a variety of interpretations. But it is pretty clear that one has not been *introspectively* aware of one's sensory states during the period of distraction (only 'pretty clear' since one could imagine that there has been rapid memory loss of such introspective awareness as one drives, but this seems to me highly unlikely as it tends to imply, or at least suggest, that all consciousness is, or is associated with, introspective consciousness, and that just seems wrong). But are distracted drivers unconscious of the sensory information which must be operating within their cognitive systems? Not according to Tye. The sensory states which represent the configuration of the road ahead as well as the orientation of the steering wheel (and many other relevant facts) seem to be doing their usual job of affecting the (short term and unremembered) beliefs one has about where and how to drive, given the standing desire not to exit the roadway, crash and burn.

So these sensory states are not merely poised to affect belief, they are *actively* affecting beliefs as the driver negotiates the roadway. Tye goes on to say that the sensory representations provide input to higher level cognitive systems 'whose job it is to produce beliefs (or desires) directly from the appropriate nonconceptual representations, *if* attention is properly focussed and the appropriate concepts are possessed' (1995, p. 138). Once again, while it is possible to read this as suggesting that attention is to be fixed upon the sensory states, this possibility ought to be resisted, since it equates consciousness with introspection.

It is not easy to think of an example of *merely* poised sensory content. Perhaps an example that Tye uses in another context will do. This is the example of Mary (p. 14ff.) who is so distracted by her thoughts that she does not notice the rose placed before her (even though she has never so much as seen a *colour* before). No beliefs or desires are formed on the basis of the sensory representations which, Tye asserts, are brought about by the rose. Although, yet again, Tye's discussion proceeds in terms of the nature of Mary's sensory experiences being available to *introspection*, let us continue to take it that it is not required for a state to be conscious that it be available for introspective awareness. Otherwise, the distinction between Tye's theory and Carruthers's dispositionalist higher order *thought* theory of consciousness would seem simply to collapse. Tye's theory can of course allow that conscious sensory content is normally available to introspective awareness in creatures, such as adult human beings, with the rather complex conceptual apparatus necessary for engaging in introspection, but his theory maintains a principled and I think essential distinction between introspective awareness and awareness *tout court*.

So it is *not* the case that it is the connection between sensory states and introspective beliefs that makes the sensory states conscious states. We may then ask if there are other restrictions on the sorts of beliefs necessary for sensory states to be conscious states. It seems that there are. Suppose there is an ANIC state S - a state, that is, which meets Tye's conditions save for that of being poised such that whenever S occurs I

directly am caused to have the (occurrent) belief that the activity level of my insular cortex has increased by 10% and let it be that S carries nonconceptual content about the level of activity of my insular cortex so that these beliefs are generally true. But I have no *sensation* of my insular cortex becoming more active; I just suddenly *think* that it is more active. This is not a case of phenomenal consciousness (although it is certainly a state of consciousness, an instance of conscious thought). This kind of direct link between sensory representation and belief is not the right sort to generate phenomenal consciousness. Real world examples of something like this imaginary case are not hard to come by. Many experiments have shown that sensory stimuli that are presented for too short a time for one to be conscious of them nonetheless have cognitive effects (see for example Kunst-Wilson, W. R., & Zajonc, R. B. (1980) or Murphy, S. T., & Zajonc, R. B. (1993)). In the latter experiment Chinese ideographs were presented to non-Chinese reading subjects who were to decide whether the ideograph represented a 'good' or 'bad' concept. Before presentation of the ideograph a human face was presented for an exceedingly short time of 4 msec. The face was either angry or happy. The expressed affect of the face influenced the subject's beliefs about the ideograph, but without consciousness of the faces. One might perhaps object that the subjects weren't forming real beliefs here, but this objection assumes too stringent a condition upon the cognitive states produced by the sensory states. The subjects of these experiments certainly were forming opinions and otherwise engaging high level cognitive function, and obviously there are plenty of conscious experiences that provoke in us only greater or lesser probabilities about the way things are.

So what are the *right* beliefs (or cognitive states)? Mere access to the belief/desire system does not by itself transform nonconceptual content into phenomenal consciousness. The obvious answer is that the sensory representations must be apt to produce beliefs about the perceptible qualities which they themselves represent. If I'm looking at a horse it will not be sufficient for phenomenal consciousness that my sensory representation of the horse induce in me the belief that there is a horse before me but rather, at minimum, a belief of the form 'a horse that looks like .... is before me'. But what fills in the dots here? No *description* will do since I can come to have beliefs about horses that meet such and such descriptions without having any states of phenomenal consciousness. We know that you could just tell me the description of how a horse that's in front of me looks and convince me it is correct (I could be blindfolded) and I could thereby come to believe that there was a horse in front of me meeting that description without enjoying any visual consciousness. Furthermore, nothing prevents a mechanism which directly instills such description based beliefs in me when I am in the presence of horses so the condition that sensory content be a 'direct influence' on one's cognitive system will not solve this problem. The fact that I could acquire this belief 'directly' from my sensory system - as in the imaginary and actual examples above - goes nowhere towards showing that phenomenal consciousness of a horse will attend the creation of the belief.

If you will forgive me, the point can be illustrated by an anecdote of a curious experience I had last summer. I was in an extremely quiet place overlooking the sea, enjoying the stillness and the few faint sounds of wind, rustling branches, birds etc. when it suddenly struck me - and this is the only way I can describe this uncanny

experience - that I was *soon going to hear something*. I was not conscious of any new sound but had the rather uneasy feeling that a sound was 'coming'. Sure enough, a short while after I could hear the distant droning of a small airplane making its way up the coast. This is an example of what has to be a state with content which is poised, in the sense of being able to 'directly' influence my cognitive system, but which was not a phenomenally conscious state. Perhaps it is possible to deny this, and to assert instead that this was a case of conscious aural experience which was 'merely' unavailable to introspection. This seems very implausible to me. There seems to be a viable distinction between a kind of pre-conscious deployment of content within our cognitive systems which is not recognised by Tye's theory unless we impose a distinction *within* the field of poised content. Thus there is a serious question about exactly what kind of poised content is 'appropriate' for the generation of consciousness. I would suggest that if this problem is genuine, it is at bottom another instance of the hard-problem or the explanatory gap.

Tye goes to a lot of trouble (ch. 2) in an effort to show that the gap is merely an illusion. His diagnosis is basically that there can be no demand for an explanation for why an identity claim holds; if x = y it is senseless to ask *why* x is identical to y. But for this to work we need a good candidate for the *physical* states that are to be identified with states of consciousness. The problem of the last paragraph reveals that poised (+ ANI) content is *not* this candidate since only *some* of this content actually underlies states of consciousness. The explanatory gap surely arises again when we're forced to limit ourselves to the entirely trivial remark that conscious states are identical to those physical states that are identical to conscious states but we have no way to telling what the relevant physical difference is between the conscious and non-conscious candidates for the identity. It always remains possible, I suppose, that we will eventually discover the candidate that maps perfectly onto conscious experience but even in that case there are constraints on what would count as *explaining* consciousness. To take a ridiculous case, suppose only PANIC states whose canonical English expression has a prime number of vowels generate states of consciousness. Although we could claim to have, we might really have, discovered the correct identity claim between physical and conscious states, this would hardly bridge the explanatory gap.

So let us return to the question of how to give a non-trivial characterization of the contents of the cognitive states necessary for the occurrence of conscious experience. Given the foregoing, the only content that might seem able to do the job here is the nonconceptual content of the sensory representation itself. One fairly natural proposal would be that only those PANIC states tending to produce beliefs (or other cognitive states) *about* those very states's contents suffice for conscious experience. But, as we've seen above, this proposal just collapses the difference between Tye's theory and the (dispositional) HOT theory of consciousness. Perhaps that is, in the end, simply what Tye's theory amounts to. This would be of some interest given that Tye and Carruthers seem to think that their theories are quite distinct.

A worse problem looms though. Beliefs cannot, by their very nature, contain nonconceptual content (although of course they can represent such content). Tye is quite clear on this restriction, for he says that conscious experiences 'arise at the

interface of the nonconceptual and conceptual domains' (p. 62). And, as we've seen, *merely* representing the nonconceptual content will not guarantee phenomenal consciousness. Call this the 'uptake problem': the problem of taking up the nonconceptual content of the sensory representations into the conceptual states of belief. It might seem that Tye introduces the notion of *phenomenal concepts* to solve this problem. But this does not actually seem to be the case. Phenomenal concepts are introduced as 'the concepts utilized when a person introspects his or her phenomenal state and forms a conception of what it is like for him or her at that time' (p. 25). A conception of what is like for a subject is a concept deployed in introspection, not a concept deployed in mere conscious awareness of the world around one. One does not need to have phenomenal concepts to have conscious perception. Tye is very liberal about allowing that animals, including some insects, have phenomenal consciousness, but he surely does not wish to grant them any introspective abilities or beliefs about what it is like to be them. Therefore such creatures will not possess any phenomenal concepts but there will nonetheless *be* something it is like to be one of them - they will enjoy phenomenal consciousness. But why? As we've seen, the fact that their nonconceptual sensory states can affect these creatures' belief/desire system does not entail that they are phenomenally conscious.

A quick and dirty answer to the uptake problem is simply to deny that there is any uptake of nonconceptual content into belief. Rather, what matters is that the linkage between the sensory contents and the belief/desire system is 'of the right kind to generate conscious experience'. One could talk here, though rather vaguely and unsatisfactorily, of there being 'enough' and 'sufficiently intimate' interaction between 'sufficiently complex', ongoing and continually changing sensory contents and similarly changing beliefs and desires. Tye's discussion of blindsight is instructive here. In short, the reason why blindsighters are not phenomenally visually conscious is supposed to be that 'there is no complete, unified representation of the visual field, the content of which is poised to make a direct difference in *beliefs*' (p. 63). The general question to ask here is *why* completeness and unity matter to phenomenal consciousness? There are neurological syndromes which seem to involve incomplete and disunified visual fields that do not destroy phenomenal consciousness. One example is that of visual neglect in which the subject is only conscious of half of the visual field. This is a case of an incomplete and disunified visual field without loss of visual phenomenal consciousness. It is easy to reply that in this case there is a complete and unified representation of *half* the visual field (which shifts around continuously with the attention of the subject). But then why don't blindsighters have complete and unified representations of 'small bits' of their visual fields, and why wouldn't this be sufficient for phenomenal consciousness? There is an obvious worry of circularity lurking here: the definition of the relevant senses of 'unified' and 'complete' can only be given in terms of consciousness itself. The point is that talk of 'complete and unified' representations of visual fields disguises a mystery: what turns some representations into conscious experiences while other representations fail to generate consciousness? This question is, of course, the good old explanatory gap reappearing again. *Some* kinds of poised content are sufficient for phenomenal consciousness but others are not. Why?

The problems discussed thus far centre on doubts about whether poised content is *sufficient* for phenomenal consciousness. It is also unclear whether poised content is *necessary* for phenomenal consciousness. The problem here is that the activation of phenomenal concepts would seem to be enough to generate a kind of phenomenal experience in the absence of any PANIC. At least, this is so if it is possible for phenomenal concepts to be active or 'applied' in the absence of what they properly apply to, a possibility which is undeniable for concepts in general, and may be a constituent feature of the *concept* of a concept, namely that they support an appearance/reality distinction. Normally, phenomenal concepts are applied to the non-conceptual contents of certain representations of either the world or the body. For example, in pain there is a non-conceptual representation of a part of the body as having certain features. Knowing what pain feels like requires the deployment of the phenomenal concept of pain (or the phenomenal concept of pain's *feel*) in the categorization of these non-conceptual contents. Presumably, Tye would endorse some tale of the neurological realization of concept application and thus there seems to be nothing preventing the 'misapplication' of phenomenal concepts if this neural realization should occur inappropriately, that is, in the absence of the appropriate non-conceptual content. Of course, such misapplication would be decidedly abnormal, but we can imagine a science fictional philosophical thought experiment in which the neurological structure of phenomenal concepts is so well understood that a machine can be constructed which can directly 'turn on' the realization of an application of a phenomenal concept, in the complete absence of any appropriate non-conceptual content. Imagine that the machine works such that when we ask someone to introspect the state of their left arm, say, it activates a particular phenomenal concept for extreme pain in the left arm. There is no reason to believe that *every* mechanism which can activate the neural machinery underlying the application of a phenomenal concept must involve a non-conceptual representation of what that phenomenal concept represents. After all, it is obvious that the signals from our imaginary machine do *not* normally covary with a state of a damaged left arm. They in fact are supposed to be nothing but a kind of list of which neural structures within the conceptual machinery are to be turned on and which turned off so as to produce the 'neural activation vector' which corresponds to the application of that phenomenal concept, along with a device which actually turns these neural structures off or on.

This is important enough to belabour. According to Tye's theory, in a normal case of introspective knowledge of what pain feels like we have two components: the non-conceptual representation of the pain (the PANIC) and the application of the phenomenal concept of pain which yields knowledge. It is explicitly granted that we can have the former without the latter. It seems equally obvious that it is physically possible that the structures that subserve application of phenomenal concept of pain could be activated in the absence of any appropriate PANIC. The imaginary machine we have envisaged brings about this latter state.

But the important question is whether there is anything it is like to be in this peculiar state. The unfortunate victims of our machine will sincerely believe that they are experiencing excruciating agony, and they will be able to describe the 'pain' in detail (since phenomenal concepts have tremendous 'fineness of grain', specificity of modality,

etc.). There seems to be a clear sense in which they are indeed feeling pain. In fact, such people would seem to meet Tye's definition of 'suffering' (see above) and it is hard to imagine genuine suffering without some kind of phenomenal consciousness.

So it appears that phenomenal consciousness can exist without there being any poised content and that poised content can exist without there being phenomenal consciousness. Without taking anything away from the interest, fruitfulness and indeed the elegance of Tye's representational theory, there still seems to be a kind of explanatory gap between representation and consciousness.

References

Block, Ned (1990). 'Inverted Earth', in J. Tomberlin (ed.) *Philosophical Perspectives*, vol. 4, pp. 53-79.

Carruthers, Peter (2000). *Phenomenal Consciousness*, Cambridge: Cambridge University Press.

Dretske, Fred (1995). *Naturalizing the Mind*, Cambridge, MA: MIT Press.

Foss, Jeffrey (2000). *Science and the Riddle of Consciousness: A Solution*, Boston: Kluwer Academic Publishers.

Kunst-Wilson, W. R. & Zajonc, R. B. (1980). 'Affective discrimination of stimuli that cannot be recognized', in *Science*, 207, pp. 557-558.

Lycan, William (1996). *Consciousness and Experience*, Cambridge, MA: MIT Press.

Murphy, S. T. & Zajonc, R. B. (1993). 'Affect, cognition, and awareness: Affective priming with optimal and suboptimal stimulus exposures', in the *Journal of Personality and Social Psychology*, 64, pp. 723-739.

Rosenthal, David (1986). 'Two Concepts of Consciousness,' *Philosophical Studies*, 49, pp. 329-59.

Seager, William (2000). 'Introspection and the Elementary Acts of Mind', *Dialogue*, 39, 1, pp. 53-76.

Tye, Michael (1995). *Ten Problems of Consciousness*, Cambridge, MA: MIT Press.

Tye, Michael (2000). *Consciousness, Color and Content*, Cambridge, MA: MIT Press.

On the Virtue of Being Poised - Reply to Seager

by Michael Tye

Department of Philosophy
The University of Texas at Austin
Austin, Texas 78750
USA

1) Seager's main focus in his probing discussion is the condition of being poised in my PANIC theory. I claim that a necessary condition upon a representational content's being a phenomenal character is that be poised. Seager's primary aim is to show that this is not the case. Poised content, he claims, can exist without phenomenal consciousness (even supposing my other conditions are met). Further, in his view, phenomenal consciousness can exist without there being any poised content. So, poised content is not sufficient either.

2) Before I turn to these points, let me mention quickly another point that Seager makes with which I am in full agreement. This is that the representational content of pain is partly 'evaluative'. Pain is not just a sensory state. It also has an affective dimension. Pain is experienced as *bad* for one.[1] It is precisely because of this that people have the cognitive reactions to them they do, reactions such as desiring to stop the pain

Of course, talk of the experienced badness of pain may sound cognitive. But it need not be understood in this way. It seems to me that the most plausible view here is that we are hard-wired to experience pain as bad for us from an extremely early age.

Consider the other side of the coin for a moment. A child as young as two months, upon tasting a little chocolate, typically behaves in a way that signifies that it wants more. The child will open and close its lips, push forward towards the chocolate, look happy. Why? The answer is that the chocolate tastes good. *That's* why the child wants more. The child's gustatory experience represents a certain taste and the child experiences that taste as good. The taste is experienced as good by the child in that the child undergoes an overall experience which represents the presence of the taste in the mouth and represents it as good.

---

[1] This is something I develop further in "Another Look at Representationalism and Pain," *Philosophical Issues*, forthcoming. It does not come out clearly enough in my *Ten Problems of Consciousness*, 1995. But see my "Blindsight, Orgasm, and Representational Overlap," *Behavioral and Brain Sciences*, Vol. 18, 1995. See also my "Orgasms Again," *Philosophical Issues*, Vol. 7, ed. by E. Villenueva, 1996.

© *A Field Guide to Philosophy of Mind* (Winter 2002)
Symposium on Michael Tye's *Consciousness, Color and Content* : 65-69.
<http://host.uniroma3.it/progetti/kant/field/tyesymp.pdf>

Intuitively, this is not a cognitive experience. It does not require concepts. It is preconceptual. For another example, consider orgasm. Orgasm is a bodily sensation, but it is not only that. The most natural description of an orgasm, and indeed of any pleasant sensation is "It feels good." One's orgasm represents a certain change in the region of the genitals as good for one, as something apt to benefit, not to harm one.[2] That isn't a conceptual response. One cannot help but feel the relevant bodily disturbance except as good. One is hard-wired by nature to experience it in this way. It is not difficult to fathom why.

3) Seager uses this point to argue that creatures without the capacity to introspect their pains still suffer. Here I am inclined to disagree. Consider the case of a pain of which you are unaware — a pain you simply fail to notice. Suppose, for example, that the pain is very brief and that, as it occurs, you are watching intently your favorite TV program. Did you *suffer* at all in undergoing that pain? Surely, intuitively, you didn't. You were blind to the pain. You had no idea at any time the TV was on that pain was present. How could you have suffered? Admittedly, your pain represented something (tissue damage, let us suppose) as bad for you. And if no error was involved, something bad did briefly happen to you. But still, it seems to me, you didn't suffer. Perhaps, this disagreement with Seager is verbal, as he says it may be. In any event, I shall not pursue it further.

4) Seager's characterization of my poisedness condition as a dispositional condition is accurate. His first objection is directed against this feature of the account. Consider a subject S_1 whose PANIC states do not, in fact, produce higher level cognitive responses, though they are disposed to do so. Now consider another subject S_2, "who is identical to S_1 save that S_2 has been modified by the attaching of a device that would block the relevant disposition — that is, make the content unavailable to higher level cognition, but *only* for those states which, in fact, are *not* going to affect higher level cognition." S_1 and S_2 are neurally and behaviorally identical, but S_2 has an inert disposition blocking device attached. On my view, according to Seager, "S_2 will have quite different states of consciousness compared to S_1." And that is very implausible.

5) That is indeed implausible. But, on my theory, properly understood, S_1 and S_2 have the same states of consciousness. Let me explain. Take a brittle wine glass. The glass is disposed to break easily, but it is treated with care and never breaks. A second wine glass, which is also treated with care and never breaks, is identical to the first save that a support has been placed inside the glass that stops it from deforming. This support would prevent the glass from breaking were it dropped, since the glass is so constituted that shattering occurs only if certain bonds in the glass break and those bonds cannot break unless some external force deforms the glass sufficiently. The second wine glass is brittle, intuitively, just as the first is. Each of them is disposed to break easily. The difference is merely that the brittleness of the second glass is *masked* by the extrinsic

[2] The suggestion that, for orgasms, goodness is part of their representational content is made in my "Blindsight, Orgasm, and Representational Overlap," *ibid.*

support. Each wine glass has the relevant disposition, since each is so constituted that, oddities aside, it would break (and thereby manifest the disposition) in the relevant circumstances (for example, dropping on hard ground).[3]

Given an understanding of this sort of dispositions, S_1 and S_2, on my account, undergo the same phenomenal states. The inert device attached to S_2 masks the disposition in S_2's case. But the disposition is still there.

6) Seager is right to say that introspective accessibility is not a blanket requirement, on my theory, for a state to be phenomenally conscious. I allow that some creatures lacking the capacity to introspect still undergo phenomenal states. Such creatures are blind to their experiences. Nonetheless, in creatures like humans, introspective accessibility in normal cases is a necessary condition. Here again Seager presents my view correctly. His next examples are ones supposedly in which there are PANIC states but these states have no phenomenal 'feel'.

The first of these is a case in which I am looking at a horse. Seager comments:

... it will not be sufficient for phenomenal consciousness that my sensory representation of the horse induce in me the belief that there is a horse before me.

I agree. As explained in my reply to Alex Byrne, what is required is that the relevant sensory representation, S, be apt for the production (in the right way) of a belief about (some of) the perceptible qualities represented by S -- qualities that I take the horse to have. Seager says, in effect, that a condition of this sort isn't sufficient either. But all that his remarks here indicate, it seems to me, is that one can *conceive* of a creature meeting my PANIC condition and yet not undergoing phenomenal states. And that is not damaging to my view; for, as I have emphasized in many places, the PANIC theory is not proposed as a conceptual truth.

7) Seager next tries to illustrate his point about an absence of sufficiency by an actual example in which he himself had (he claims) the "rather uneasy feeling that a sound was 'coming' " even though at the time he had no auditory experience (all was quiet, we are told). A little later, he heard the distant noise on an airplane. He remarks:

This is an example of what has to be a state with a content which is poised,... but which was not a phenomenally conscious state.

However, by his own characterization, the state *was* phenomenally conscious. For he describes it as an uneasy *feeling*. What kind of feeling was it? That is difficult to say without further information. Perhaps the feeling was, at least in part, of very slight

[3] This view of dispositions and the example just offered are in Johnston, M. 1992 "How To Speak of the Colors," *Philosophical Studies*, 68, pp. 221-263.

bodily vibration. Alternatively, perhaps Seager merely had the expectation or occurrent thought that a sound was coming, and perhaps this cognitive state had accompanying it a linguistic auditory image (so that it seemed to Seager that he was uttering internally "A sound is coming" with his usual accent and stress). Granted, this would involve a kind of auditory experience and Seager says that all was quiet. But perhaps what he means is that all was quiet *externally*, as represented by his experiences.

8) The next objection Seager raises is of a general sort. He thinks that there is no satisfactory way of explaining how the nonconceptual content of experiences is 'taken up' into the conceptual content of beliefs. He calls this "the uptake problem." I am not sure exactly what is troubling Seager here. Consider the case of perceptual experiences. Beliefs about the directly accessible, perceptible qualities of things are based upon perceptual experiences representing those qualities. If one thinks of the basing relation which connects the relevant experiences with the relevant beliefs in internalist terms, then Seager is right: there is indeed a problem. For the nonconceptual content of a perceptual experience cannot form a premise for reasoning to a conceptual conclusion. Reasoning or inference requires conceptual premises. But that is not how I think of the basing relation. The relation, in my view, is an externalist one. There are reliable mechanisms connecting the contents of sensory states and the relevant belief contents, mechanisms whose functioning underwrites the transition from something looking *F*, for example, to the belief that it is *F*, given that one has the concept *F*.

9) Seager's final objection involves a case in which, he asserts, phenomenal consciousness exists without any poised content. Suppose that a given individual is *not* in the relevant PANIC state for pain, but the individual is in a (cognitive) state *S* that misapplies the phenomenal concept PAIN. In these circumstances, according to Seager, there is something it is like for the individual in undergoing *S*. For, Seager claims, the individual believes that he is in pain and he is able to describe further features of the 'pain' he takes himself to be undergoing. In short, "[t]here seems to be a clear sense in which [this individual] is feeling pain."

10) I beg to differ. By hypothesis, Seager's example involves a *mis*application of the phenomenal concept PAIN. How could this be if the individual were indeed feeling pain? Indeed, so far as Seager has described the case, it could be one of absent qualia.

Phenomenal concepts, in my view, are exercised (in the first person case) in our *awareness* of our phenomenal states via introspection. They enable us to become aware of the felt character of our phenomenal states. As noted earlier, without such concepts, we would be 'blind' to our 'feels'. We would be in much the same state as the distracted driver who is thinking hard about philosophy, say, as he drives along the highway.[4] The driver is unaware of how the road ahead looks to him, of the visual experiences he is

[4] This case is due to David Armstrong. See his *A Materialist Theory of Mind*, London: Routledge and Kegan Paul, 1968.

undergoing; for his attention is focused elsewhere. But the experiences are there alright. He still *sees* the road in front of him. How else does he keep the car on the road?

Cognitive awareness of our own feelings itself feels no special way at all. Phenomenal character attaches to experiences and feelings (including images), and not, I maintain, to our cognitive responses to them.[5] Admittedly, as I noted earlier, phenomenal concepts are concepts that *dispose* their possessors to form images or phenomenal memories of the relevant experiences (among other things); but the concepts themselves do not have an experiential character. This being so, there is nothing it is like for the individual of Seager's example to be in state *S*. She certainly believes that pain is present, and she certainly takes herself to be suffering pain. In reality, she isn't. There is no pain for her to suffer.

[5] As noted earlier, I do not deny that some cognitive responses have associated linguistic, auditory images. I should add that in saying here that phenomenal character attaches to experiences and feelings, of course, I do not mean to commit myself to the view that phenomenal character is a quality of experiences and feelings.