

Discussion

CONSCIOUSNESS AND NONCONCEPTUAL CONTENT*

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, in particular the use it makes of “nonconceptual content”.¹



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, p. 136, note omitted; cf. 1995, p. 101)²

“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 (p. 138).

It seems to me that Tye himself has supplied compelling counter-examples against this proposal, namely various perceptual illusions, in particular the Müller-Lyer illusion (1995, p. 102; 2000, p. 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.³

2. THE PANIC THEORY

The PANIC theory is this: “phenomenal character is one and the same as Poised, Abstract, Nonconceptual, Intentional Content” (2000, p. 63; cf. 1995, p. 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, p. 138; 2000, p. 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, p. 62; cf. 1995, p. 138). A visual experience as of a tomato is poised, because it typically causes a belief about the tomato "if attention is properly focused" (p. 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, p. 62). We can speak derivatively of poised contents: a content is poised iff it is the content of some poised state.

Finally, 'nonconceptual'. 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, p. 139). A *state* is nonconceptual iff it has nonconceptual content.

According the PANIC theory, an intentional state lacks phenomenal character just in case it isn't poised, or doesn't have abstract or nonconceptual content. Letting X be the elusive ingredient that makes an intentional state one with phenomenal character, we can summarize the PANIC account of this ingredient with the slogan 'X = P + A + N'.

If $X = P + A + N$, then the significance of this discovery can hardly be exaggerated. The next section examines some problems with N.

3. NONCONCEPTUAL CONTENT

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 2 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, p. 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 ... F⁴ So, for example, if a subject believes that cranberries are red, or that cranberries are not red, or that everything red is colored, then she possesses the concept *red*. And if she possesses the concept *red* then she has *some* belief whose content can be specified using the English word ‘red’.

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

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, p. 104).⁵ 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*.⁶

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 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.⁷ (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 “protopro-

positional content”. These abstract objects are built to Russellian specifications: a protopositional content *is* a simple sort of Russellian proposition, while a scenario content is something more complicated, but likewise constructed from materials at the level of reference (Peacocke, 1992, ch. 3). The contents of belief and thought, on the other hand, are exclusively conceptual.⁸

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.⁹ 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 red₁₉, for example, is phenomenally different from my experience of red₂₁, even though I have no stored memory representations of these specific hues and hence no such concepts as the concepts *red₂₁* and *red₁₉*. These points generalize to the other senses. Phenomenal character, and hence phenomenal content, on my view, is nonconceptual (1995, p. 139; cf. 2000, pp. 61–62).

That is, to possess the concept *F* (i.e. to believe that . . . *F* . . .) one must have, at least, “the stored memory representation *F*”. And because it is possible to have a visual experience of red₂₁, without having “the stored memory representation *red₂₁*”, one does not have to possess the concept *red₂₁* in order to have that visual experience. Therefore, a visual experience of red₂₁ 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 red₁₉, where the proposition that the tomato is red₁₉ is the very same kind of proposition – a Fregean Thought, perhaps – that she can believe.¹⁰

Tye’s official argument for nonconceptual content establishes, at best, the *state* view. But Tye in fact holds the *content* view.¹¹ 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. (This is argued for in Byrne, 2001a, section 6.1.)

The PANIC theory, then, implies the content view: beliefs (thoughts, judgments) have conceptual content – content that is not composed of “concepts” or Fregean senses – and perceptions have content of another kind. (Examples of nonconceptual content are Russellian, Lewis-Stalnakerian, and Peacockean (scenario) contents.)

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, p. 243).¹² And, indeed, proponents of the content view invariably endorse the claim that linguistic content is also Fregean.¹³

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.¹⁴ Neither has he given any positive characterization of nonconceptual content. Because of this, it is completely obscure why nonconceptual content (on the content conception) is part of ingredient X.

4. X = P + A + N REVISITED

Although Tye may have misidentified ingredient X, 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 argued in Byrne, 2001a, 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.

I shall now outline an argument for this claim, based on Jackson’s (1982) knowledge argument together with a perceptive remark of Lewis’s.¹⁵ Assume, first, that knowing what it’s like to enjoy an experience is propositional knowledge.¹⁶ When black-and-white Mary sees a ripe tomato for the first time, and thereby comes to know what it’s like to see red, she comes to know some proposition. If one were forced to choose a sentence to express this proposition, a plausible candidate would be ‘Seeing red is like *this*’, where we imagine Mary uttering this sentence while looking at a tomato. So, assuming for the moment that the proposition Mary learns is linguistically expressible, we may write it thus:

(M₁) Seeing red is like *this*.

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

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

For an intentionalist like Tye, Mary comes to know M₂, not by directly introspecting her *experience*, but by attending to the colors in the scene before her eyes: “Our attention goes outside in the visual case, for example, not to the experience *inside* our heads. We attend to one thing – the external surface and qualities – and yet *thereby* we are aware of something else, the ‘feel’ of our experience” (2000, pp. 51–52).¹⁷ In other words, Mary is in a position to know M₂ once she knows:

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

Note that M₃ 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 M₃ 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, p. 281). Therefore, since knowing M₃ would help imprisoned Mary to know what it’s like, the proposition M₃ cannot be taught by a *lesson*.

But what is a “lesson”? In one sense, showing Mary a ripe tomato is giving her a lesson, but obviously that is not what Lewis means. Instead, it’s clear that he means *linguistic* lessons. No matter how many books imprisoned Mary reads, and lectures she hears, she won’t come to know what it’s like to see red. And this is not because there are some sentences that Mary can’t understand. Although she hasn’t had the experience of seeing red objects, that does not prevent her from understanding any linguistic expression (so, for example, she can understand the word ‘red’ while imprisoned). Of course, there will be uses of demonstratives that could not occur in lessons Mary has while imprisoned, in particular an utterance of ‘An experience as of red objects represents them like *this*’ in the presence of a tomato. And such an utterance of that sentence expresses –

we have been supposing – the proposition M_3 . But this does not mean that the proposition M_3 – if it really is expressed by that sentence – could not be taught to imprisoned Mary. Plausibly, any proposition expressed using a demonstrative could be expressed in a demonstrative-free way: for example, the proposition expressed by ‘That man is drinking a martini’ (pointing at Tye) is arguably expressed by the demonstrative-free sentence ‘Tye is drinking a martini’. Assume this is correct. Then, if M_3 really is expressed by an appropriate utterance of ‘An experience that represents objects as red represents them like *this*’, we could teach M_3 to imprisoned Mary: no demonstration of ripe tomatoes is needed.

All the premises are now in place (albeit with minimal defense). If M_3 can be linguistically expressed, then Mary can know M_3 while imprisoned, and thereby know what seeing red is like. But she can’t know this while imprisoned. Therefore M_3 can’t be linguistically expressed. Our supposition that M_3 is expressible using a demonstrative is a ladder that must be kicked away: in using a demonstrative, we were trying to say what can’t be said. We can, however, *communicate* or *convey* M_3 , by uttering the sentence ‘An experience that represents objects as red represents them like *this*’ in the presence of a ripe tomato; at least, M_3 can be communicated in this way to those who have the appropriate sort of experience. (And, I presume, I have succeeded in communicating M_3 to *you*.) For familiar Gricean reasons, a proposition can be communicated by uttering a sentence in a context, even if the proposition is not the semantic content of that sentence relative to that context. Hence, it doesn’t follow from the fact that M_3 can be communicated by uttering a sentence in a context, that M_3 is the semantic content of that sentence relative to that context; neither does it follow that M_3 is the semantic content of *some* sentence.¹⁸

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.

The provisional conclusion, then, 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.

NOTES

* Many thanks to Michael Glanzberg, Ned Hall, Sally Haslanger, Jim John, Sarah McGrath, Jim Pryor, Mark Richard, Susanna Siegel, Robert Stalnaker, Ralph Wedgwood, and Steve Yablo. The present paper is an excerpt from Byrne, 2001a.

¹ 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 Tye later adds a complication (2000, pp. 139–140) in the style of Fodor's asymmetric dependency account (Fodor, 1990, ch. 4); this is not relevant here.

³ 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, p. 119). Dretske discusses this problem in an endnote (174, n. 13), and gives a version of the reply mentioned above.

⁴ See 1995, p. 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.

⁵ See also 1995, p. 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).

⁶ Cf. Crane, 1992, p. 143.

⁷ Strictly speaking, (b) should be: the content of perception is *at least partly* nonconceptual. (See, e.g., Peacocke, 1992, p. 88.) This complication will be ignored.

⁸ 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.

⁹ The point to follow is an elaboration of Byrne, 1996, p. 264, n. 6. Because the richness argument at best supports the state view, Byrne (1996, pp. 263–264) 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, p. 155).

¹⁰ The richness argument is in embryo form in Evans, 1982, p. 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, pp. 67–68; 1998; for a more guarded view of the argument, see 2001b) and Heck (2001, pp. 489–490); Heck’s version of the richness argument is discussed in Byrne, 2001a. (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, pp. 56–60; 1998) on the ground that demonstratives like ‘that shade’ can capture the content of color experience (see also Brewer, 1999, pp. 170–174; Kelly, 2001). However, McDowell appears to concede that the richness argument provides a *prima facie* consideration in favor of the content view.

¹¹ He confirmed this in correspondence. (For a slight complication – not examined further here – see note 13 below.)

¹² 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).

¹³ 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, p. 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, p. 695; cf. 2000, p. 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).

¹⁴ According to Tye, the representational vehicles of experiences have a “topographic or maplike structure” (1995, p. 121; cf. 2000, pp. 70–74), unlike the representational vehicles of beliefs, which have a sentencelike structure (1995, p. 100) (so Tye thinks there is a language of thought, although not a language

of experience). One might try to argue from these differences in *representational vehicles* to a difference in the kinds of *contents represented*. However, Tye does not supply any such argument, and there is no indication that he thinks one could be supplied. Moreover, it would be a confusion to think that a difference in representational vehicles *entailed* a difference in contents represented.

For an examination of other arguments for the content view, see Byrne, 2001a.

¹⁵ For another way of approaching the same conclusion, see Byrne forthcoming; an important related discussion is in Thau, 2002.

¹⁶ 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, p. 16).

¹⁷ See also Byrne, 2001b; Dretske, 1995, ch. 2; Shoemaker, 1994.

¹⁸ 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

- Brewer, B. (1999): *Perception and Reason*, Oxford: Oxford University Press.
- Byrne, A. (1996): ‘Spin Control: Comment on John McDowell’s *Mind and World*’, in E. Villanueva (ed.), *Philosophical Issues*, vol. 7, Atascadero, CA: Ridgeview.
- Byrne, A. (2001a): DON’T PANIC: Tye’s intentionalist theory of consciousness. (Available at: mit.edu/abyrne/www/DON'TPANIC.pdf.)
- Byrne, A. (2001b): ‘Intentionalism Defended’, *Philosophical Review* 110, 199–240.
- Byrne, A. (forthcoming): ‘Something about Mary’, *Grazer Philosophische Studien* (special issue on the philosophy of Terry Horgan).
- Chalmers, D.J. (1996): *The Conscious Mind*, Oxford: Oxford University Press.
- Crane, T. (1992): ‘The Nonconceptual Content of Experience’, in T. Crane (ed.), *The Contents of Experience*, Cambridge: Cambridge University Press.
- DeLucia, P.R. and Hochberg, J. (1991): ‘Geometrical Illusions in Solid Objects Under Ordinary Viewing Conditions’, *Perception and Psychophysics* 50, 547–554.
- 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.

- Jackson, F. (1982): 'Epiphenomenal Qualia', *Philosophical Quarterly* 32, 127–136.
- Kelly, S.D. (2001): 'The Non-conceptual Content of Perceptual Experience: Situation Dependence and Fineness of Grain', *Philosophy and Phenomenological Research* 62, 601–608.
- 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, 403–431.
- Peacocke, C. (1992): *A Study of Concepts*, Cambridge, MA: MIT Press.
- Peacocke, C. (1998): 'Nonconceptual Content Defended', *Philosophy and Phenomenological Research* 58, 381–388.
- Peacocke, C. (2001a): 'Does Perception Have a Nonconceptual Content?', *Journal of Philosophy* 98, 239–264.
- Peacocke, C. (2001b): 'Phenomenology and Nonconceptual Content', *Philosophy and Phenomenological Research* 62, 609–615.
- Shoemaker, S. (1994): 'Self-knowledge and "Inner Sense". Lecture III: The Phenomenal Character of Experience', *Philosophy and Phenomenological Research* 54, 219–314.
- Stalnaker, R. (1998a): 'What Might Nonconceptual Content Be?', in E. Villanueva (ed.), *Philosophical Issues*, vol. 9, Atascadero, CA: Ridgeview.
- Stalnaker, R. (1998b): 'Replies to Comments', in E. Villanueva (ed.), *Philosophical Issues*, vol. 9, Atascadero, CA: Ridgeview.
- Stanley, J. and Williamson, T. (2001): 'Knowing How', *Journal of Philosophy* 98, 411–444.
- 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, 695–697.

Department of Linguistics and Philosophy
Massachusetts Institute of Technology

Alex Byrne