

Review

Consciousness, Color and Content

by Michael Tye.

Cambridge, Massachusetts: The MIT Press. pp. xiii + 198.

PAUL NOORDHOF

Consciousness, Color and Content is a defence of the representationalist approach to phenomenal consciousness (hereafter *representationalism*) by one of its most prominent advocates. The book is clearly written, attractively laid out, robust in argument and, almost unerringly, it focuses on the substantial issues central to current debate rather than getting side-tracked. Tye does the invaluable job of adding to our understanding of the range of options we may take up in understanding the nature of phenomenal consciousness as well as providing a snapshot of contemporary philosophical debate on those issues. He also displays a sensitivity to the relevance of empirical work which will be welcome to readers of *Mind & Language*.

In a short review, I must be selective. The book ends with a piece of applied philosophy of mind. Given that Tye's view of the nature of consciousness is correct, which creatures are likely to be conscious? The book begins with a discussion of Tye's favoured approach to Frank Jackson's famous knowledge argument against physicalism. I shall pass over these parts. The heart of the book is part II which contains Tye's defence of representationalism. In the first section of this review, I will raise a methodological concern about this defence which has consequences for its successful development. The second chapter of the book contains Tye's treatment of the explanatory gap problem: the fact that it seems that there will always be an explanatory gap between what we are told by theories of the nature of the brain and environment and the nature of phenomenal consciousness. Tye claims that this chapter is self-standing. I think it is in some tension with other parts of Tye's approach and this will be matter I discuss in the second section of the review.

1. Tye's defence of representationalism

Phenomenal consciousness is consciousness of what it is like to be in a certain mental state (hereafter, the phenomenal content of the mental state).

Address for correspondence: University of Nottingham, University Park, Nottingham NG7 2RD, UK.

E-mail: paul.noordhof@nottingham.ac.uk

Mind & Language, Vol. 18 No. 5 November 2003, pp. 538–545.

© Blackwell Publishing Ltd. 2003, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

Representationalists claim that the phenomenal content of mental states is determined by their representational properties. As Tye puts it, necessarily, experiences alike in their representational contents are alike in their phenomenal content (p. 45). Tye puts forward a particular account of the nature of phenomenal content. It is a species of poised nonconceptual, representational content (p. 60). They are poised because they are ready and available to have direct impact on beliefs and desires (p. 62).

One of the primary motivations for this thesis is the intuition of transparency. When we have a visual experience of a wine glass sitting on a table, then a proper characterisation of what the experience is like is given in terms of the wine glass sitting on the table. Our introspection of what our experience is like reveals no further properties, it just goes straight through, as it were, to the properties of the wine glass. It is in this sense that the experience is transparent. According to Tye, the best explanation of these facts is that the representational properties of the experience of the wineglass—the properties of the experience in virtue of which it is an experience of a wineglass—determines the phenomenal content of the experience.

Put like this, representationalism about phenomenal content is very plausible. It is true that there is a question-mark over an appeal to this argument given the existence of theories of perception which deny that the phenomenal content of perceptual experience is fixed by representational properties while emphasising that perception involves a certain kind of non-representational openness to the world. Nevertheless, Tye might fairly wonder whether such theories can deny that perceptual states have representational properties in the way that he would seek to claim that they do. So I shall put these theories aside.

It is worth noting that representationalism is not discredited by the fact that a proper characterisation of the phenomenal content requires mention of something non-representational, namely a wine glass. The crucial point is that representational properties determine the phenomenal content of the experience. This is quite compatible with what is represented being non-representational.

Once this is recognised, there seems to be a strengthened argument for representationalism which serves to bring the methodological issue into focus. Suppose that the opponent insists that there is a non-representational feature of experience which contributes to what the experience is like. If the opponent's grounds for insistence is that the experience itself presents an additional feature of the world not mentioned so far in the characterisation of the representational content of the experience, then we just have a case in which the representational content of the experience has been incompletely characterised. We need to mention this additional feature and, along with that, attribute a further representational property to the experience. On the other hand, if the presence of the non-representational feature is not revealed by the presentation of some feature of the world, then we need to consider the way in which its presence is revealed. Presumably, the answer is that it is revealed in an experience of the experience which possesses it. A higher-order experience involves a presentation of the non-representational feature of the lower-order experience. In which case, the higher-order experience should be

attributed the corresponding representational property involved in the presentation of the non-representational feature of the lower-order experience. Either way, we have no counterexample to representationalism.

The argument just given means that representationalism may be defended merely by pointing out presentational differences in some experience or other corresponding to any phenomenal difference mentioned. For the most part, Tye implicitly adopts this strategy. For every case, he points out presentational differences and takes these to indicate that there are representational differences (pp. 69–97). However, a puzzling feature of his approach is that he is also committed to a certain view about the nature of representational properties. Yet he does not really consider whether the apparent counterexamples might present a challenge to his account of representational properties or the legitimacy of the presentations he identifies.

Tye holds that:

S represents that P = df. (a) If optimal conditions were to obtain, S would be tokened in *c* [a creature] if and only if P were the case; moreover, in these circumstances, S would be tokened in *c* because P is the case (b) If there is some other feature Q which covaries with P, then were P to fail to causally covary with Q, Q would not satisfy (a) (with Q in place of P) but P would still satisfy (a) (pp. 136–140).

In the case of evolved creatures, optimal conditions obtain if a subject's perceptual mechanisms are functioning as they are designed to function. There must be no interference in their functioning (p. 138). Tye justifies this account in terms of the verdicts it gives regarding inverted earth cases (where sky is yellow, sun is blue, etc.). If subjects from earth have colour-inverting lenses installed in their eyes to compensate for the colours on inverted earth, it would be appropriate to count them as continuing to see the sky as blue because the optimal conditions would be determined by what the sensory apparatus producing S were designed to indicate by S on earth. In the case of non-evolved creatures otherwise identical to us—the proverbial swamp people—with compensating lenses, Tye argues that they will misperceive colours both on earth and on inverted earth. If a normal subject without lenses would see something blue on earth, the swamp twin would see it as yellow with the lenses and vice versa on twin earth. Because these subjects misrepresent, the conditions cannot be optimal (pp. 138–139).

In passing, let me note that it is not clear that this argument works. Tye needs to show that there was a sense in which one could apply the notion of optimality to the non-evolved creatures. Only then could their experience represent something one way or another. However, Tye concludes that the swamp person would see blue objects as yellow given inverting lenses in order to conclude that the conditions are not optimal. This suggests that he is prepared to assign phenomenal content independent of the account given above.

Be that as it may, I am more interested in the methodological issue. The Müller-Lyer illusion provides an illustration in point. I assume that there will be occasions

when a subject looks at the figures in optimal conditions. There is nothing interfering with the functioning of his or her perceptual organs. Nevertheless, the lines are presented to have different lengths. This suggests that Tye's account of representational properties is incorrect. I envisage that Tye might say that the circumstances aren't optimal during illusions. The question is what notion of optimality is in play. It cannot be just that something gives rise to misrepresentation.

Consider a case Tye explicitly covers in the book: after-images (pp. 83–86). What is the thing with which a subject's experience is causally correlated in optimal circumstances in order to have an after-image? It is hard to think of anything which fits the bill. After-images don't exist and it is not as if there is an object which would be presented in the same way as an afterimage to play the role instead. One way that Tye could deal with this is to argue that the account of representational properties given above only applies to primitive perceptual elements. Notice, though, that, in both cases, we have no reason to think that the representational picture is false. Indeed, it seems to give the right account of the phenomenal content of perceptual experience. It is just that this treatment is not well supported by the particular account of representational properties Tye endorses.

Tye's apparent insensitivity to this possibility is revealed in his treatment of the intuition that different people might see colour in different ways as a result of differences in their sensory constitution. We could capture this thought by saying that the representations of colours play a different cognitive role due to differences in the capacity to distinguish between colours. The cognitive architecture of one subject represents two colours as quite distinct whereas that of another does not. Although the second subject may be more likely to make mistakes, it does not follow that either subject misrepresents the way colours are. By contrast, Tye seems forced to claim that one subject would always misrepresent colours in such circumstances (pp. 89–93, 104–109). Instead, we can say that the colours are represented as having different properties due to the difference in cognitive role they play. This would involve revising one's account of representational properties. I should note that, at one point, Tye allows that different classifications in play with regard to the degree of orange and red in a colour can give rise to phenomenal and, thereby, representational differences. However, he does not explain how this touches his account of representational properties (pp. 92–93).

Tye's commitment to a certain account of representational properties means that he cannot simply argue that there are apparent representational differences when his opponent argues that an experience has a phenomenal content which is partly non-representational. He must offer a back-up proof that what is represented does stand to experience in the appropriate way for his account of representational properties to work. This he does not do. By the same token, he undermines the force of his position by his commitment to a dubious account of the representational character of perception.

2. The Explanatory Gap and Representationalism

The problem of the explanatory gap can be posed by contrasting the following. It is an a priori truth that:

Solidity = The disposition to retain shape and volume (or something like this).

If we consider the properties whose joint instantiation, together with the particulars which possess them, constitutes the fact that certain molecules are not free to move around, we can, a priori, appreciate that the property of being solid is instantiated. By contrast, there seems no a priori truth of the form:

Phenomenal pain = the F (where 'phenomenal pain' picks out the phenomenal property constitutive of the feeling of pain and 'the F' is a description in terms of physical properties) (pp. 21–23, 33–35).

Tye allows that physical properties include higher order properties variably realised by the properties of micro-physics (say). This enables him to state the problem just in terms of property identity. Tye suggests that the reason why we feel that there is an explanatory gap in the second type of case stems from the nature of our phenomenal concepts.

There appear to be two components to Tye's account and it is not entirely clear how they relate. The first component of his proposal seems to go like this. When we consider an identity such as phenomenal pain = the F, we feel that it is mistaken because our concept of a physical state is not apt to trigger in us imaginings and memories of pain in the way that our concept of the phenomenal state is. As a result it appears that the physical state fails to do justice to the phenomenology. But this is just an artifact of the way the two concepts function (pp. 27–28).

This proposal doesn't seem to work. Tye takes it to apply not only to identity statements involving brain states but also to those involving functional states (p. 23, fn. 3). However, in this case, our concept of the relevant functional state, F, would also seem apt to trigger imaginings and memories. The only functional state that might plausibly be identified with phenomenal pain is one whose defining functional role includes imaginings and memories of the kind appropriate to capture the role of pain in learning and avoidance. If Tye's explanation of the intuition that the phenomenology is not captured were correct, we should feel that the phenomenology *is* captured with regard to the functional identity. Yet we do not.

I suspect that Tye places most weight on the second component of his approach. He suggests that the key feature of phenomenal concepts is that they are directly recognitional. They are not applied to experiences via intermediate reference-fixing descriptions (pp. 28–29). It is this fact which gives rise to the illusion of an explanatory gap.

Tye compares the situation with the case of water. There is no explanatory gap between H₂O and water because the following deduction of facts about water from facts about H₂O is possible.

- (1) Water = The F.
- (2) H₂O = The F.
- (3) There is H₂O in place p.

Therefore,

- (4) There is water in place p.

(1) is supposed to mark the place of an a priori truth. It might be something of the sort: 'Water is the bearer of many of the following features: being a liquid, filling lakes and oceans, coming out of taps, being called 'water' by English-speaking experts, being necessary for life on the planet, falling from the sky', and so on (pp. 29–30). This description fixes the reference of the concept water in the actual world. (2) and (3) are supposed to be empirical truths. Given (2) and (3), it follows a priori that there is water in place p. By contrast, we cannot deduce that phenomenal pain is instantiated because there is no equivalent a priori truth of the form:

- (1') Phenomenal P = the F (or an F).

This might seem straightforwardly false. Isn't phenomenal pain, to take one example, *the property which is instantiated as a result of tissue damage and whose instantiation brings about pain behaviour*? It is at this point that Tye's use of the phrase 'phenomenal. . . .' becomes important. He does not need to deny that it is a priori true that the property of being in pain is *the property which is instantiated as a result of tissue damage and whose instantiation brings about pain behaviour*. However, it does not follow that *phenomenal pain* is the unique property which satisfies this description. Other phenomenal properties may also satisfy it. Different people may feel pain in different ways. Since these would be all phenomenal properties of pain, we might talk of phenomenal pain₁, phenomenal pain₂, and so on, with similar labelling for the corresponding phenomenal concepts (p. 27, fn. 6). Moreover, it is not a priori true that phenomenal pain fits the description since it is conceivable that Epiphenomenalism is true.

I don't think this explanation of the source of the explanatory gap works. Tye faces a dilemma. It can be brought into focus by the following question: 'Must the properties referred to in the description *The F* be narrowly physical or can they be broadly physical?' Narrowly physical properties are just those identified by a correct physics which bears some resemblance to our own. Broadly physical properties are properties which supervene upon narrowly physical properties. A brief look at the case of water suggests that Tye supposes that the properties identified need only be broadly physical. However, if that is right, we seem to have the following candidate for (1) regarding phenomenal pain.

- (1'') Phenomenal pain is *the property which triggers application of the concept of phenomenal pain on the basis of introspection when introspection is functioning normally*.

Is the italicised description of a broadly physical property? If Tye says 'no', then physicalism is false. On the other hand, if he says that it is a broadly physical property, then we have a contingent a priori truth to use in the following deduction.

- (2') X-firing is *the property which triggers application of the concept of phenomenal pain on the basis of introspection when introspection is functioning normally*.

Therefore,

- (3') X-firing is phenomenal pain.

(2') would be an empirical truth given to us by a correct theory of introspection in just the same way as we discover that H₂O has the properties mentioned in the description of water. It seems that given this, we can conclude, a priori, that X-firing is phenomenal pain.

It might be questioned whether (1'') is really a contingent a priori truth but by Tye's own lights it is. According to Tye, it is an a priori truth that phenomenal concepts are involved in the direct recognition of phenomenal states. This is partly cashed out as the idea that:

- (I) In normal conditions, introspection of the phenomenal property P causes the subject to apply the concept of phenomenal P (pp. 28–29).

In which case, (1'') comes out a priori true.

It would not help to insist that the properties referred to in the description *The F* should be narrowly physical. It is very implausible that we fix the referent of the term 'water' by a description using terms referring only to narrowly physical properties (as Ned Block and Robert Stalnaker have recently pointed out). We would have to identify water as the possessor of the various micro-physical properties that realise the properties Tye identified, for instance, the microphysical properties that realise taps and the English speech community's use of the term 'water'. In that case, the putative reference-fixing properties would not be available to us a priori. We do not introduce the concept water via micro-physics. In which case, the explanatory gap would appear to open up in the case of relating water to H₂O. It is clear that it does not.

Tye's commitment to representationalism appears to make matters worse. One might think that his proposal goes wrong by focusing on any old a priori reference fixing description instead of an a priori description of the very nature of phenomenal pain (just as we described solidity as the disposition to retain shape and volume). Tye's representationalism would appear to license the following inference.

- (4) Phenomenal pain is the representation of tissue damage.
- (5) The representation of tissue damage is the relational property S would have if, if optimal conditions were to obtain, S would be tokened in c if and only if c suffered tissue damage, etc. etc.

Therefore,

- (6) Phenomenal pain is the relational property S would have if, if optimal conditions were to obtain, S would be tokened in c if and only if c suffered tissue damage, etc. etc..

According to Tye, phenomenal pain has a representational nature. I take it that this is an a priori truth. Hence (4) is a priori true. (5) would also come out a priori true by Tye's lights since, I take it, the justification that Tye offers for his account of the nature of representational properties is a priori. For the deduction to work, it need only be an empirical truth. So nothing rests on this. From (4) and (5) we can conclude (6). It would be interesting to consider whether representationalism is committed to there being no explanatory gap and what consequences this has for the plausibility of representationalism.

*Department of Philosophy
University of Nottingham*

Review

Complex Demonstratives: A Quantificational Account

by J. C. King.

Cambridge, Mass.: MIT Press, 2001. pp. ix, 206.

EMMA BORG

Complex demonstratives are phrases formed by concatenating a demonstrative term, like the English expressions ‘that’, ‘this’ or ‘those’, with some descriptive component, say a common noun or a restrictive clause. In this way we get syntactically complex expressions like ‘this girl’ or ‘that philosopher from UC Davis who likes skiing’. The question, which has been on the lips of many philosophers of language recently, then is: what is the correct semantic analysis of these syntactically complex expressions? Specifically, are they semantically complex or not? The problem here should be obvious. We have expressions which due to the presence of a demonstrative term (and thanks to the analysis of demonstratives *simpliciter* as paradigm terms of direct reference by Kaplan et al) seem to demand analysis as genuine referring terms. Yet, due to the presence of a descriptive component, these expressions also seem to warrant analysis as descriptive, quantified phrases. The ‘orthodox view’ (King’s label) takes the demonstrative feature of complex demonstratives to be fundamental and suggests an analysis of these expressions as simple referring terms. An issue for this kind of approach is then how to find room for the descriptive component in a directly referential account and the usual response has been to put it in the character.¹

Jeff King, however, as his book title suggests, rejects the orthodox view in favour of a quantificational analysis. For King, the syntactic complexity of these expressions, their containment of an overt descriptive component, and their behaviour in certain key contexts, all points to them requiring a quantificational analysis; and the kind of quantificational theory he proposes is certainly ingenious. His suggestion, in broad outline, is that ‘that’ has a lexical entry expressing a four-place relation: ‘__ and __ are uniquely __ in an object x and x is __’ (p. 43). So, in any ‘that’ phrase (as he commonly terms complex demonstratives), two argument places will be filled by the properties picked out by predicates overtly appearing in the surface syntax (i.e. the properties relating to ‘F’ and ‘G’ in ‘that F is G’),

Address for correspondence: Philosophy Department, University of Reading, Reading RG6 6AA, UK.

Email: e.g.n.borg@reading.ac.uk

¹ See, e.g., D. Braun, 1994, ‘Structured Character and Complex Demonstratives’, *Philosophical Studies* 74, 193–219; E. Borg, 2000, ‘Complex Demonstratives’, *Philosophical Studies* 97, 229–249.

while two—the second and third places above—will be filled by properties determined by the intentions of the speaker (pp. 43–4). So, what are these contextually determined properties, and why have two of them?

The first such property relates to the kind of contact an agent has with the object she wants to talk about, namely, is it an object of (past or present) perception, or is it something she knows about only via description? In the former case, imagine a speaker can see a person, *b*, that she wants to talk about and she utters ‘That man is tall’. Here the speaker’s intention is perceptually based, she intends to talk about the thing she can see, and hence the property determined will simply be *is identical to b*. Thus we can partially complete the relation expressed by her utterance, as follows: ‘*is a man* and *is identical to b* are uniquely—in an object *x* and *x is tall*’. However, sometimes the agent will be thinking of the object purely descriptively and her intention will then be to talk about ‘whatever satisfies some property or conjunction of properties *O**’ (p. 44). So, imagine that someone is thinking of a man simply as ‘the best contemporary basketball player’, then an utterance of ‘That man is tall’ will yield a different partially completed relation: ‘*is a man* and *is the best contemporary basketball player* are uniquely—in an object *x* and *x is tall*’ (p. 34).

This still leaves open another slot to be contextually filled and King reserves this place for one or other of two properties of pairs (or collections of) properties: the final slot is filled either by the property of *being jointly instantiated in w, t* (where ‘*w, t*’ gives the world and time of the utterance) or by the property of *being jointly instantiated*. This ‘higher order’ argument slot is important because the rigidity of the ‘that’ phrase will depend on which property is in play here. If the property determined by the speaker’s intention is the property of *being jointly instantiated in w, t*, then the expression will be rigid (in all circumstances of evaluation, the object we are concerned with will be just that object which jointly instantiates the properties in the first two argument places *in w, t*; i.e. in the world and time of the context of utterance). Whereas, if the property in this slot is the property of *being jointly instantiated*, then the expression will be non-rigid (in any circumstance of evaluation we will be interested in the object which jointly instantiates the other properties *at that circumstance of evaluation*). King’s provisional suggestion (p. 84) is that the two contextual slots in ‘that’ phrases are not filled independently of one another: a property determined by a perceptual intention in the first slot will lead to the property of *being jointly instantiated in w, t* in the second slot, while a descriptively determined property in the first slot pairs with the property of *joint instantiation* in the second. Thus, in our above two examples, we have two distinct completed relations. In the perceptual case the completed relation is: ‘*is a man* and *is identical to b* are uniquely *jointly instantiated in w, t* in an object *x* and *x is tall*’. Since the instantiation relation is relativised to a world we are guaranteed to pick out the same object at all circumstances of evaluation, thus capturing the rigidity of perceptual uses of ‘that’ phrases. In the descriptive case, however, the completed relation is: ‘*is a man* and *is the best contemporary basketball player* are uniquely *jointly instantiated* in an object *x* and *x is tall*’, where the instantiation relation is not relativised to a world and thus where a different object may be picked out at

different circumstances of evaluation, so long as the selected object is a man and is the best contemporary basketball player at that world. Though the official line of the book is that the two slots for properties determined by speaker intentions are not filled independently, we should note that King does not rule out the possibility of 'autonomous' slot filling in a final account. Thus, ultimately, up to four combinations of our two contextually determined properties might occur (pp. 80–4).

Specifying a lexical entry for 'that' (at least when it appears concatenated with descriptive material) which is 'gappy' and where (at least some of) these gaps are to be filled by contextually determined properties makes 'that', uniquely according to King, a context-sensitive determiner. Unlike 'every', 'some' and 'the', with which it semantically groups, the meaning of 'that' requires contextual supplementation, via the intentions of the speaker, before it can deliver anything truth-evaluable. King does not rule out another, potentially similar, kind of context sensitivity which does permeate all quantified noun phrases, namely the appeal to a context of utterance to establish the domain of quantification. However, he stresses that accounts of domain constraint which trace this context sensitivity to the underlying syntax of the expressions in play tend to locate it in the common nouns themselves; whereas, on King's account of complex demonstratives, we have a genuinely context-sensitive determiner.²

King's account is complicated, and I'm suppressing many important details here, but one question which should spring immediately to mind for anyone familiar with the issues surrounding demonstratives is: why does King want an account which allows some expressions of the form 'that F' to be non-rigid? Kripke and Kaplan showed us that demonstratives are rigid, so why should we think complex demonstratives are any different? It is in addressing this question that I think King's book is at its strongest, for he marshals an extremely impressive range of cases which seem to show that, on some occasions at least, 'that' phrases behave like quantified noun phrases, i.e. they have narrow scope, non-rigid readings. In short, in many kinds of contexts complex demonstratives do behave just as we would expect if they belonged to the semantic category of quantified phrases. King's arguments here range from the semantic (e.g. intuitions about when certain utterances containing 'that' phrases are true or false, and when indirect speech or propositional attitude attributions are correct) to the purely syntactic (e.g. concerns about movement and weak crossover effects) and, though the advocate of a direct reference analysis of complex demonstratives may have a response to make to King's examples on a case-by-case basis, there is no doubt that, taken together, the data adds up to a powerful case for a quantificational approach.

Now, since I've advocated a directly referential analysis in the past myself, I'd obviously like to think that King's cases do not provide knock-down evidence against the orthodox view, and that, on a case-by-case basis, there do remain avenues for a defender of direct reference for complex demonstratives to explore. Here isn't the

² See King chapter 4; J. Stanley and Z. Szabo, 2000, 'On Quantifier Domain Restriction', *Mind and Language* 15: 219–261.

place to undertake this kind of detailed defence, but I would like to mention just one issue. For with respect to some of the semantic evidence, it seems to me that we do need to be careful to disentangle data to be explained by a correct account of 'that' phrases and data to be explained by a correct account of indirect discourse or propositional attitude attributions. For it's not entirely clear that all of King's examples reveal features of the meaning of 'that' phrases, as opposed to highlighting what many have thought is a common discrepancy between literal semantic analyses and features affecting judgements of what is communicated in a given context.

Take King's objection (pp. 115–6) that a direct reference account cannot deal with the intuitive truth of a report of the form 'A said that that woman is beautiful', where the original speaker did indeed say 'that woman is beautiful' but intended to refer (unbeknownst to the original speaker only) to a man dressed as a woman. In the first case, I think we might question the judgement of intuitive truth for the report, or at least question its semantic import. As an act of *oratio recta*, the report is certainly correct, and it also seems true if the term 'woman' is being mentioned rather than used in the report, but, given that all bar the speaker are aware of the male status of the subject, an appropriate act of *oratio obliqua* should, one might think, not repeat the original speaker's mistake but should instead make sure that reference to the right object is secured (where the best way to achieve this is to utilise a predicate all know to hold of the object). Secondly, it seems that our reporter could (at least sometimes) report A by uttering 'A said that that man is beautiful', while pointing at the person concerned, and still intuitively, I think, say something true. Yet in this case, if acceptable reports can be made using incompatible predicates, then the suggestion that one or other of the reports is informative as to the semantic status of the original utterance seems to be undermined.³ The thought here is that we need a further argument, from some intuitions about when an act of indirect discourse is true or false, to a claim about the semantic value of the original utterance since, in most cases, it's features outside the purely semantic

³ Jeff King (personal communication) has pointed out that the above fact might be held to *support* his position, since he can allow that both reports are literally correct ('A said that that woman is beautiful' is true when 'that woman' takes narrow scope with respect to 'A said that', while 'A said that that man is beautiful' is true when the 'that' phrase takes wide scope). Thus we both agree that both reports can be true, though I take this fact to show that indirect speech reports are sensitive to non-semantic features, while King takes it to support a semantic analysis permitting scope ambiguities for 'that' phrases. Clearly this is a point which bears further discussion, but I might just note one point which *perhaps* tells in favour of my interpretation. Imagine the following scenario: there is a creature, *o*, which A believes is a cat, while everyone else agrees that it is a dog. However *o* is in fact a large weasel. A says 'That cat is fat', which, on the above model, might be correctly reported either as 'A said that that cat is fat' or 'A said that that dog is fat' (where the latter report picks out the object in a way the audience believes is correct). Yet, if the latter report is correct, this is not so easily explained by giving the 'that' phrase wide scope, since the object in question is not a dog but a weasel. I'd suggest that the simplest explanation of this case is that indirect speech reports are sensitive to pragmatic features of both (original) context of utterance and the context of report, so that intuitions about speech reports do not map neatly to semantic claims; but, as noted above, this point bears further discussion.

that affect indirect discourse. In many cases in the book, where we are asked to make assessments about the correctness or otherwise of an indirect discourse or a propositional attitude attribution, I felt that some additional argument was needed on King's behalf to show us that the intuitions in play tell us about the semantic properties of 'that' rather than the pragmatic properties of communication. Drawing apart these kinds of features won't help in all cases (for instance, some other manoeuvre will certainly need to be deployed in the face of 'quantifying in' uses, like 'every student admires that professor who first taught her logic'), but, together with other refinements (like a clearer identification of the actual explanandum for a theory of complex demonstratives), there may still be scope for hanging on to direct reference for all demonstratives, simple or complex.⁴

Whether this scope is real or imaginary, however, my real worry with King's account surfaces at a rather more brute level than the kind of close argument and complex response his data certainly demands. For, on his positive account, speakers are thought to learn a four-place relation when they learn the meaning of 'that' (or at least 'that F') and they are taken to have some very sophisticated intentions concerning the filling of the two contextually determined argument places in this four-place relation. This semantic complexity is, at least *prima facie*, pretty remarkable, and one might question the psychological plausibility of supposing that the nascent language user, on adding 'that F' to her vocabulary, comes to have intentions regarding such relations as identity, or that on hearing 'that F' she is sensitive to an incredibly complex pair of intentions held by her interlocutor. It seems surprising, to say the least, to think of the young child, withdrawing her hand on hearing her mother's utterance of 'that dog bites', as revealing a competence to grasp, and reason about, another person's intentions concerning properties of collections of properties, like co-instantiation across possible worlds. Furthermore, it's a competence which, King is perfectly willing to allow, may *not* be required for utterances of 'that' *simpliciter* (e.g. if bare demonstratives are terms of direct reference, while 'that' phrases are quantified expressions). Yet it seems to me unlikely that we will find the kind of step change in the conceptual abilities of nascent speakers envisaged here in moving from understanding bare demonstratives to understanding complex ones.

Complex demonstratives become, on King's account, incredibly sophisticated pieces of linguistic machinery, requiring some extremely refined intentions on the behalf of interlocutors. Noting this complexity isn't, of course, anything like a knock-down argument against King's view—after all, lots of linguistic theories posit more complexity in their semantics than can be found at the surface level, and the idea that kids 'cognize' more than they can consciously access is commonplace. Yet it remains to be seen whether King can really relegate the complexity to be found in his account to the merely tacit (for these properties fill in the lexical entry

⁴ Concerning the syntactic evidence King canvasses, see also K. Johnson and E. Lepore (2002), 'Does Syntax Reveal Semantics?', *Philosophical Perspectives*, and King's (unpublished ms, available on King's website) response 'Syntactic Evidence for Semantic Claims'.

for 'that' and the competent speaker is usually held to have conscious access to the meanings of her own words, and they provide part of what the speaker intends to say by her utterance), and, if not, whether the complex intentions posited can fit with such constraints as learnability. Furthermore, it's simply not obvious that we are inclined to take these complex intentions to impact on grasp of meaning in the way the current account predicts (e.g. see n. 30, p. 185). Say I utter 'That F is G' with the two property places determined by speaker intentions filled by a property determined by a perceptual intention and the property of being jointly instantiated at *w*, *t*, but my audience take the relevant properties to be, say, one determined by (an appropriate) descriptive intention plus the property of being jointly instantiated at *w*, *t*.⁵ In this case, according to the current theory, you fail to understand my utterance, yet since we both arrive at the same object, both predicate F and G of it, and both take that very (actual world) object to be relevant in all other circumstances of evaluation, it's certainly not obvious that such a failure will ever be judged, by ordinary interlocutors, to have taken place (certainly, it's hard to envisage a situation where the addressee's 'mistake' will ever come to light). Or again, imagine we are discussing works of art and I point to a Titian and say 'That artist had a fantastic palette'. Since neither of us have had (past or present) perceptual contact with Titian, it seems that we must be thinking of the artist descriptively, but how on earth are you to know if I'm thinking of him as the painter of that picture, or the Italian artist who painted 'Venus and Mars', or as my favourite Renaissance artist? And, more to the point, why should you care? It seems to me that, so long as you understand every word in the syntax of my utterance, and you understand the way those words are put together, then you're going to understand the meaning of my utterance, regardless of your (in)ability to second guess the way I'm thinking about the painter. Here, as elsewhere in accounts of literal linguistic meaning, I think we are wise to fight shy of moves to incorporate speaker intentions into the semantic analyses of noun phrases.

Despite the above reservations, however, it's clear that King's book is a model of its kind. Essentially, King sets out to do two things (and he does them both very well): he sets out the argument for a quantificational analysis of complex demonstratives and he proposes, and powerfully argues for, a novel kind of quantificational theory. Ultimately, whether the reader endorses one or both (or neither) of these two points is somewhat beside the point here, for in the kind of systematic, exhaustive marshalling of data and the close, careful argument to be found here, King's book sets the standard for future accounts of complex demonstratives on either side of the divide.

*Philosophy Department
University of Reading*

⁵ A combination disallowed by the official line of the book, but explicitly left open as a possibility in a final correct account; see pp. 81–4.